

Department of Mathematics, Royal Holloway College, University of London
July 2016

Small Salem Graphs

PhD thesis by Jonathan Charles James Cooley
Supervised by Professor James F McKee

Declaration of Authorship

I Jonathan Charles James Cooley hereby declare that this thesis and the work presented in it is entirely my own. Where I have consulted the work of others, this is always clearly stated.

July 2016

Abstract

The aim for this thesis was to produce the first systematic catalogue of small Salem graphs, and to illustrate and enumerate those with interesting properties. This is done in the central section, chapters 3 to 5. That is preceded by two introductory chapters, the first dealing with definitions and motivation, the second concerning the computational methods used to construct the catalogue. A new isomorphism testing algorithm is presented which has proved highly successful in practice, but an example is constructed in which it is useless. A complete classification of circulant Salem graphs is conjectured. For Salem graphs, their Salem number and their Mahler measure are precisely the same thing, and some observations were made which enabled the completion of the classification of all graphs with Mahler measure below ϕ , the very well-known ‘golden number’. The final chapter is an exposition of the paper which reports this completion of the classification.

Contents

Chapter 1 Definitions and motivation

- 1.1 Salem numbers 1.2 Totally real algebraic integers, and graphs
- 1.3 Linking the previous two sections 1.4 Combinatorial objects and polynomials
- 1.5 Mahler measure 1.6 Structure and contribution of this thesis

Chapter 2 Computation

- 2.1 Acknowledgement 2.2 Jacobi's algorithm
- 2.3 PARI: advantages and problems 2.4 An isomorphism algorithm
- 2.5 Some Salem graphs with unusual distance tables
- 2.6 The structure of two "almost isomorphic" graphs

Chapter 3 A catalogue of small Salem graphs

- 3.1 Building the catalogue 3.2 Numerical results
- 3.3 The 16 cyclotomic graphs with 6 vertices or fewer
- 3.4 The 3 Salem graphs with 4 vertices 3.5 The 16 Salem graphs with 5 vertices
- 3.6 The 74 Salem graphs with 6 vertices

Chapter 4 Minimal Salem graphs

- 4.1 The 18 GRINs and their computational verification
- 4.2 When GRINs can coexist in SCSGs 4.3 Minimal non-Salem graphs

Chapter 5 Other small Salem graphs of special interest

- 6.1 Complementary graphs 6.2 Cospectral graphs
- 6.3 Integral graphs 6.4 Regular graphs

Chapter 6 Non-bipartite graphs of small Mahler measure

- 6.1 Introduction 6.2 Theorem 6.3 All kites have small Mahler measure
- 6.4 All balloons have small Mahler measure 6.5 Details of some computations
- 6.6 Some graphs that do not have small Mahler measure
- 6.7 All large enough, connected, non-cyclotomic, non-bipartite graphs of small Mahler measure are either kites or balloons
- 6.8 Proof of corollary 6.2.1 6.9 Final remarks

List of tables

Table 3.1 Numbers of cyclotomic and Salem graphs with a given number of vertices

Table 5.1 Numbers of complementary graphs with fewer than twelve vertices

Table 5.2 Numbers of sets of isospectral Salem graphs with fewer than twelve vertices

Table 5.3 Numbers of integral Salem graphs and trivial Salem graphs with fewer than twelve vertices, listed according to their number of edges

Table 5.4 Numbers of circulant Salem graphs with fewer than 41 vertices, arranged by their n -regularity

Table 5.5 Numbers of regular Salem graphs with fewer than twelve vertices arranged by their n -regularity

Table 6.1 Mahler measures of small kites and balloons

Table 6.2 Small connected graphs with Mahler measure strictly between 1 and ϕ

Chapter 1

Definitions and motivation

An abstract of this thesis is on page 3. After the main definitions, §1.6 describes the structure and contribution in more detail. Each chapter begins with a paragraph on its original content.

In this chapter only the speculation in §1.3.2 and the order and layout of the content are original.

1.1 Salem numbers

1.1.1 Cyclotomic polynomials

Cyclotomic polynomials have been very extensively studied; it is clear that mathematicians find them interesting.

The n th cyclotomic polynomial Φ_n is

$$\Phi_n(x) = \prod_{\substack{1 \leq k \leq n \\ \gcd(k,n)=1}} \left(x - e^{\frac{2\pi i k}{n}} \right)$$

where $e^{\frac{2\pi i k}{n}}$ are the n th roots of unity. The restriction $\gcd(k,n)=1$ picks out those roots which are primitive and also makes the degree of Φ_n equal to $\phi(n)$ where ϕ (here) is Euler's totient function.

The properties of cyclotomic polynomials are remarkable, and their appearance is deceptively simple. It is clear from the definition that Φ_n is monic. Less obviously it has coefficients which are integers and which are palindromic, and it is irreducible over \mathbb{Q} . It is also of even degree, and the first 104 cyclotomic polynomials (where n has no more than two odd prime factors) have all their coefficients drawn from the set $\{0, \pm 1\}$ [W1].

The only two (important) exceptions are:

$$\Phi_1 = x-1 \quad \text{and} \quad \Phi_2 = x+1.$$

Clearly these do not have even degree and the former is not even palindromic, but they are easy to identify and to remove when they are factors of polynomials of higher degree, and in the present work they only occur as repeated factors, which are both even and palindromic (§1.3.1).

Φ_n is sometimes *defined* as the minimum polynomial (over \mathbb{Q}) of any primitive n th root of unity.

Integral, monic, palindromic polynomials of even degree are closed under multiplication, but not under factorization:

$$x^4 - 3x^2 + 1 = (x^2 + x - 1)(x^2 - x - 1).$$

1.1.2 “Almost cyclotomic” polynomials

Interesting new mathematics often arises when definitions are relaxed slightly and the results are investigated, so it is perhaps natural to look for polynomials which have all their roots but one on the unit circle. We now show, however, that this is impossible; “almost cyclotomic” polynomials must have two (closely-related) roots (at least) off the unit circle.

Let $f \in \mathbb{Z}[x]$ be monic and irreducible, with $\deg(f) \neq 1$ and with a root z_1 on the unit circle. Then the coefficients of f are palindromic and the roots of f consist entirely of mutually reciprocal pairs.

$f(\bar{z}_1) = 0$ because the coefficients are real and $\bar{z}_1 = \frac{1}{z_1}$ because z_1 is on the unit circle.

Let

$$g(x) = x^{\deg(f)} f\left(\frac{1}{x}\right).$$

The coefficients of g are the coefficients of f in reverse order, and in particular the constant term of g must be one because f is monic.

$$g(z_1) = z_1^{\deg(f)} f\left(\frac{1}{z_1}\right) = z_1^{\deg(f)} f(\bar{z}_1) = 0$$

so f and g share a root. This means that their gcd must be non-trivial, but f is irreducible so $g = cf$ where c is a constant. Further, c must be ± 1 because of the constant term of g .

$f(1) = g(1)$ because the sum of the coefficients is the same. But if $f = -g$ then $f(1) = 0$ so f has $x-1$ as a factor. But f is irreducible, so $f = x-1$ which is of no interest. So $f = g$ and

$$f(x) = x^{\deg(f)} f\left(\frac{1}{x}\right).$$

This has palindromic coefficients and is called a reciprocal polynomial.

The constant term is one, so zero is not a root, and

$$f(x) = 0 \Leftrightarrow f\left(\frac{1}{x}\right) = 0.$$

Since we have excluded the trivial cases $x = \pm 1$, or $f = x \mp 1$, we have that x and $\frac{1}{x}$ are distinct. So all the roots of our “almost cyclotomic” polynomial come in mutually reciprocal pairs, where the two members of each pair are distinct, so the degree of f is even.

1.1.3 Salem numbers

We have seen that if $f \in \mathbb{Z}[x]$ is monic and irreducible with $\deg(f) \neq 1$ and with a root on the unit circle, then it is a reciprocal polynomial with roots in mutually reciprocal distinct pairs. So it cannot have a single real root.

The closest that an “almost cyclotomic” polynomial can get to being cyclotomic is to have two real roots, with the same sign, one with modulus less than one and the other with modulus greater than one. If the signs are negative we prefer the polynomial $f(-x)$, or we simply change the signs of the real roots. Then the root which is greater than one is called a **Salem number** and the other is its reciprocal.

1.1.4 Polynomials which are not irreducible

Let $f \in \mathbb{Z}[x]$, as before, be a reciprocal polynomial with exactly one pair of real roots (excluding ± 1) and all its other roots (at least one pair) on the unit circle. What can we say if f is not irreducible?

The two real roots must be roots of the same irreducible factor, and the other irreducible factors have all their roots on the unit circle and so must be cyclotomic, by a theorem of Kronecker [M]. If the irreducible

factor with the two real roots has degree greater than two, then we still have a Salem number and its reciprocal, though we don't know the degree of the Salem number without going to the trouble of removing the cyclotomic factors.

The insidious case, which does not give a Salem number, is when the two real roots are the only roots of a quadratic polynomial. Such a quadratic must be monic, and the constant term must be one because the roots multiply to one. So it must be of the form

$$x^2 - ax + 1$$

where the coefficient of x must be negative to make the roots positive, and the positive integer a must be greater than two to make the roots distinct and real.

A number like the larger root (which is outside the unit circle but with "all" its Galois conjugates inside) is in fact called a Pisot number.

1.2 Totally real algebraic integers, and graphs

1.2.1 Characteristic polynomials

A root of a polynomial which has integral coefficients is called an algebraic number, and if the polynomial is monic it is called an algebraic integer. If the polynomial is irreducible it is called the minimum polynomial of the algebraic integer and the degree of the algebraic integer is said to be the same as the degree of the polynomial. The other roots of the minimum polynomial are called the (Galois) conjugates of the algebraic integer. If they are all real then the algebraic integer is called totally real.

The adjacency matrix of a graph is a very special case of a Hermitian matrix, so all its eigenvalues are real. The characteristic polynomial of a graph is clearly monic with integral coefficients (though not necessarily irreducible) so eigenvalues of adjacency matrices are totally real algebraic integers.

It is also known that all totally real algebraic integers are eigenvalues of graphs [BEG].

1.2.2 The interval $[-2, 2]$

Let θ be a totally real algebraic integer and define the interval $I_\theta = [\theta_{min}, \theta_{max}]$ where θ_{min} and θ_{max} are the least and greatest of the Galois conjugates of θ . If J is another interval then we know the following [M]:

- If the length of J is strictly less than 4 then there are finitely many θ with $I_\theta \subseteq J$.
- If the length of J is strictly greater than 4 then there are infinitely many θ with $I_\theta \subseteq J$.
- If J is $[-2, 2]$ (and hence if J is $[a, a+4]$ where $a \in \mathbb{Z}$) then

$$I_\theta \subseteq J \Leftrightarrow \theta = 2\cos(q\pi) \text{ where } q \in \mathbb{Q}.$$

In other cases where the length of J is exactly 4 nothing is known.

Clearly the interval $[-2, 2]$ is important in the study of totally real algebraic integers, as will be

1.3 Linking the previous two sections

1.3.1 A transformation

Consider the transformation

$$\mathbb{Z}[x] \rightarrow \mathbb{Z}[z] \quad f(x) \mapsto z^{\deg(f)} f\left(z + \frac{1}{z}\right).$$

We claim that this defines a one-to-two correspondence between roots x which are in the interval $[-2, 2]$ and pairs of roots z which lie on the unit circle. If we start with a polynomial which has exactly one root outside the interval, we end up with an almost cyclotomic polynomial which in turn gives a Salem number. Graphs with adjacency matrices whose characteristic polynomials satisfy this starting condition are called (non bipartite) **Salem graphs**.

Proof of claim. Let

$$x = z + \frac{1}{z},$$

$$\text{so } z^2 - xz + 1 = 0 \text{ and } z = \frac{x \pm \sqrt{x^2 - 4}}{2}. \quad (\star)$$

Then if $x = \pm 2$ we have $z = \pm 1$ (twice) and if $x \in (-2, 2)$ we have


$$|z| = \sqrt{\frac{x^2}{4} + \frac{4 - x^2}{4}} = 1.$$

It can now be seen why the factors $x+1$ and $x-1$ always occur in pairs (§1.1.1).

Conversely if z is on the unit circle, then

$$x = \operatorname{Re}\left(z + \frac{1}{z}\right) = \operatorname{Re}(z + \bar{z}) = 2\operatorname{Re}(z) \in [-2, 2].$$

1.3.2 Origin of the transformation



It is not clear who first used this transformation or how it was devised, but the following speculation is surely convincing.

Draw the interval $[-2, 2]$ directly above the unit circle (with the same width) and simply draw vertical lines to create the association between points in the interval and pairs of points on the circle.

Clearly we have

$$\operatorname{Re}(z) = x - 2 \quad \text{and} \quad \operatorname{Re}^2(z) + \operatorname{Im}^2(z) = 1.$$

This gives

$$\operatorname{Im}(z) = \sqrt{1 - \frac{x^2}{4}}$$

from which we can construct

$$z = \frac{x \pm \sqrt{x^2 - 4}}{2}.$$

So $(2z - x)^2 = x^2 - 4$ and $4z^2 - 4xz + 4 = 0$ giving

$$x = z + \frac{1}{z}.$$

It is worth pointing out that if we start with a more general interval $[-n, n]$ then we end up with

$$x = \frac{n}{2} \left(z + \frac{1}{z} \right),$$

so if $n \neq 2$ all hope of ending up with a monic polynomial is lost.

A very similar transformation is familiar from solving cubic equations. Consider the following example, chosen to avoid complex numbers and to minimize the need to use fractions:

$$x^3 - 3x^2 - 9x + 27 = 0.$$

Put $x = y + 1$ to remove the square term and get

$$y^3 - 12y + 16 = 0.$$

Now use a close relative of the current transformation, $y = z + \frac{4}{z}$ to get a quadratic in z^3 , and the rest is straightforward.

This beautiful piece of mathematics is not nearly as widely known as might be expected, presumably because it is usually necessary to find the cube root of a complex number which introduces a huge dose of tedious arithmetic.

1.3.3 Bipartite graphs

Let the vertices of a bipartite graph be ordered so that all the vertices of one vertex set (between which there are no edges) come before the vertices of the other vertex set. Then the adjacency matrix has the

block structure $\begin{pmatrix} \mathbf{0} & \mathbf{B} \\ \mathbf{B}^T & \mathbf{0} \end{pmatrix}$.

Suppose further that the matrix has an eigenvalue λ with corresponding eigenvector $\begin{pmatrix} \mathbf{e}_1 \\ \mathbf{e}_2 \end{pmatrix}$.

$$\text{So } \begin{pmatrix} \mathbf{0} & \mathbf{B} \\ \mathbf{B}^T & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{e}_1 \\ \mathbf{e}_2 \end{pmatrix} = \begin{pmatrix} \mathbf{B}\mathbf{e}_2 \\ \mathbf{B}^T\mathbf{e}_1 \end{pmatrix} = \lambda \begin{pmatrix} \mathbf{e}_1 \\ \mathbf{e}_2 \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} \mathbf{0} & \mathbf{B} \\ \mathbf{B}^T & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{e}_1 \\ -\mathbf{e}_2 \end{pmatrix} = \begin{pmatrix} -\mathbf{B}\mathbf{e}_2 \\ \mathbf{B}^T\mathbf{e}_1 \end{pmatrix} = -\lambda \begin{pmatrix} \mathbf{e}_1 \\ -\mathbf{e}_2 \end{pmatrix}.$$

This means that the eigenvalues are symmetrical about zero so we can never get exactly one eigenvalue outside the interval $[-2, 2]$ suggesting that bipartite graphs are no use for the present purpose.

However, we know from considering elementary graphs (§5.2.3) that graphs with no odd cycles, the bipartite graphs, have characteristic polynomials in which the degrees of all the terms have the same parity. Also if n and m have the same parity with $m < n$ then the degrees of all the terms in the expansion of $x^n \left(x + \frac{1}{x}\right)^m$ are even.

So instead of using the transformation

$$f(x) \mapsto z^{\deg(f)} f\left(z + \frac{1}{z}\right)$$

we can use

$$f(x) \mapsto z^{1/2 \deg(f)} f\left(\sqrt{z} + \frac{1}{\sqrt{z}}\right)$$

and proceed as before. Bipartite graphs with adjacency matrices whose characteristic polynomials have a single root greater than 2 are called (bipartite) **Salem graphs**.

1.4 Combinatorial objects and polynomials

We now have a way of generating polynomials (and hence Salem numbers) which are as close as possible to cyclotomic (without actually being cyclotomic) both from a large number of graphs which are bipartite and from a large number of those which are non-bipartite.

The idea of associating polynomials with graphs (or with other combinatorial objects) is not uncommon, but usually the interest centres on the graph, and the polynomial encapsulates information relating to it. Here on the other hand, the interest is rather in the polynomial which is generated from the graph.

The least of the known Salem numbers, Lehmer's number (1.17628...) was found by this method because its negative is the larger real root of the Alexander polynomial of a particular pretzel knot [Bo]:

$$x^{10} + x^9 - x^7 - x^6 - x^5 - x^4 - x^3 + x + 1.$$

1.5 Mahler measure

For a monic polynomial $p(z) \in \mathbb{Z}[z]$, its Mahler measure, written $M(p)$, is defined by

$$M(p) = \prod_{p(\alpha)=0} \max(1, |\alpha|),$$

where multiple roots contribute to the product according to their multiplicity.

The Mahler measure of a graph G , written $M(G)$, is defined (in [MS]) to be the Mahler measure of its associated reciprocal polynomial.

The crucial point for this thesis is that, for a Salem graph, its Salem number and its Mahler measure are exactly the same thing.

If a graph G has characteristic polynomial $\chi_G(x)$ then it will later be convenient to write the Mahler measure of the graph $M(G)$ in terms of the eigenvalues of G by using equation (★) in §1.3.1.

$$M(G) = \prod_{\chi_G(\lambda) = 0, |\lambda| > 2} \frac{|\lambda| + \sqrt{\lambda^2 - 4}}{2}.$$

1.6 Structure and contribution of this thesis

The thesis is in three sections:

Section 1 is introductory.

Chapter 1 is mainly concerned with definitions and motivation; especially how the main idea arises naturally from extending the familiar idea of cyclotomic polynomials.

Chapter 2 is about the computational system and the computational methods used to carry out the study. The one rather specialist built-in function which was used pervasively, and an original algorithm which was highly effective in practice (but far from perfect in theory) are both described in detail.

Section 2 is the first systematic study of small Salem graphs.

Chapter 3 is a catalogue of very small Salem graphs and an enumeration of slightly less small Salem graphs. This was a considerable computational achievement at the time, but will probably seem trivial before long. Detailed information on the graphs which were enumerated but mostly did not otherwise appear in the thesis will be freely available and perhaps of some future use. This hope is supported by the fact that, in their paper [GM], the authors thank the present author for independently confirming their calculations — a fairly small task which would have been huge had the catalogue not existed.

Chapter 4 looks in more detail at two types of minimal graphs.

Chapter 5 looks in more detail at some other special cases.

Section 3 is **Chapter 6**. The equivalence, for a Salem Graph, of its Salem number and its Mahler measure (§1.5) allowed observations to be made which lead to the completion of the classification of Salem graphs of ‘small’ (defined in chapter 6) Mahler measure. This section is an exposition of the paper which contains that completion.

Chapter 2

Computation

The presentation of §2.2, the observations and work-arounds of §2.3 and the whole of §2.4 to §2.6 are original.

2.1 Acknowledgement

All computations carried out for this thesis were done with PARI, a free computer algebra system primarily aimed at number theorists, maintained at the University of Bordeaux [P].

PARI programs are interpreted rather than compiled; speed is achieved by making use of the very large number of built-in objects and functions which have been optimized and compiled. Only one of these functions has both been very extensively used throughout the production of this thesis and uses an algorithm, the details of which are not widely known. This is the function `qfjacobi`, where `qf` (quadratic form) is a prefix applied to a group of PARI functions and `jacobi` refers to Jacobi's algorithm for finding the eigensystem of a real symmetric matrix.

2.2 Jacobi's algorithm

It is very well known that a non-singular matrix A can be written as $A = PDP^{-1}$ where the columns of P are the eigenvectors of A , and D is a diagonal matrix with the eigenvalues of A on its leading diagonal [W2]. Jacobi's algorithm achieves this decomposition by applying a sequence of very simple rotations (and their inverses) to A . At each stage the rotation is chosen to make the largest off-diagonal entry of A zero. It is not at all surprising that each iteration undoes some of the "good work" done by previous iterations. It is extremely far from obvious, however, that we get convergence to the desired result at a very acceptable rate.

2.2.1 Givens matrices

This simple two-dimensional rotation is used to introduce a zero into a vector:

$$\text{Given } a \text{ and } b, \text{ find } r \text{ and } \theta \text{ so that } \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} = \begin{pmatrix} r \\ 0 \end{pmatrix}.$$

Clearly $r = \sqrt{a^2 + b^2}$ because rotations preserve lengths; it is then easily verified that $\cos \theta = \frac{a}{r}$ and $\sin \theta = \frac{-b}{r}$.

A matrix doing a similar job in more than two dimensions is now¹ called a Givens matrix. Such a matrix represents a rotation about an axis through the origin which is perpendicular to two other co-ordinate axes.

¹ The algorithm was introduced in 1846 by Carl Gustav Jacob Jacobi (1804–1851). The matrices were given their present name when they became practically useful during the lifetime of Wallace Givens (1910–1993), and were made more widely known by his work.

Specifically, given i and j :

- $g_{ii} = g_{jj} = \cos \theta$ (which we write as c for brevity in what follows)
- $g_{ij} = -g_{ji} = \sin \theta$ (which similarly we write as s)
- $g_{kk} = 1 \quad \forall k \neq i, j$ and
- all other entries are zero.

So, for example, the Givens matrix in six dimensions with $i = 5$ and $j = 3$ is:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & c & 0 & -s & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & s & 0 & c & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

2.2.2 Eigenvalues

Now we consider $A' = G^T A G$ where A is real symmetric (initially an adjacency matrix) and G is a Givens matrix, noting that A' and A are similar because $G^T = G^{-1}$.

The following, with $n = 6$, $i = 5$ and $j = 3$ is large enough to see clearly what happens and enable mental extrapolation to the general case. It is open to the charge of being a typically poor undergraduate “proof by example”, but it is not hard to convince oneself that the criss-cross pattern shown by the shaded bars will always occur, and a more rigorous (and almost certainly more opaque) presentation, such as [BDKM], would probably trouble the conscientious reader to devise a pencil-and-paper example along the lines of what follows anyway.

$$\text{Let } A = \begin{bmatrix} A_{11} & A_{12} & A_{13} & A_{14} & A_{15} & A_{16} \\ A_{12} & A_{22} & A_{23} & A_{24} & A_{25} & A_{26} \\ A_{13} & A_{23} & A_{33} & A_{34} & A_{35} & A_{36} \\ A_{14} & A_{24} & A_{34} & A_{44} & A_{45} & A_{46} \\ A_{15} & A_{25} & A_{35} & A_{45} & A_{55} & A_{56} \\ A_{16} & A_{26} & A_{36} & A_{46} & A_{56} & A_{66} \end{bmatrix}.$$

The three important entries

- A_{ii} ($= A_{55}$ in this case),
- A_{jj} ($= A_{33}$ in this case) and
- $A_{ij} = A_{ji}$ ($= A_{35} = A_{53}$ in this case) all need to be considered individually.

Entries outside the shaded areas do not change.

$$A' = G^T A G = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & c & 0 & s & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -s & 0 & c & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} & A_{13} & A_{14} & A_{15} & A_{16} \\ A_{12} & A_{22} & A_{23} & A_{24} & A_{25} & A_{26} \\ A_{13} & A_{23} & A_{33} & A_{34} & A_{35} & A_{36} \\ A_{14} & A_{24} & A_{34} & A_{44} & A_{45} & A_{46} \\ A_{15} & A_{25} & A_{35} & A_{45} & A_{55} & A_{56} \\ A_{16} & A_{26} & A_{36} & A_{46} & A_{56} & A_{66} \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & c & 0 & -s & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & s & 0 & c & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

$$= \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & c & 0 & s & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & -s & 0 & c & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} A_{11} & A_{12} & cA_{13}+sA_{15} & A_{14} & cA_{15}-sA_{13} & A_{16} \\ A_{12} & A_{22} & cA_{23}+sA_{25} & A_{24} & cA_{25}-sA_{23} & A_{26} \\ A_{13} & A_{23} & cA_{33}+sA_{35} & A_{34} & cA_{35}-sA_{33} & A_{36} \\ A_{14} & A_{24} & cA_{34}+sA_{45} & A_{44} & cA_{45}-sA_{34} & A_{46} \\ A_{15} & A_{25} & cA_{35}+sA_{55} & A_{45} & cA_{55}-sA_{35} & A_{56} \\ A_{16} & A_{26} & cA_{36}+sA_{56} & A_{46} & cA_{56}-sA_{36} & A_{66} \end{bmatrix}$$

$$= \begin{bmatrix} A_{11} & A_{12} & cA_{13}+sA_{15} & A_{14} & cA_{15}-sA_{13} & A_{16} \\ A_{12} & A_{22} & cA_{23}+sA_{25} & A_{24} & cA_{25}-sA_{23} & A_{26} \\ cA_{13}+sA_{15} & cA_{23}+sA_{25} & s^2A_{55}+2scA_{35}+c^2A_{33} & cA_{34}+sA_{45} & (c^2-s^2)A_{35}+sc(A_{55}-A_{33}) & cA_{36}+sA_{56} \\ A_{14} & A_{24} & cA_{34}+sA_{45} & A_{44} & cA_{45}-sA_{34} & A_{46} \\ cA_{15}-sA_{13} & cA_{25}-sA_{23} & (c^2-s^2)A_{35}+sc(A_{55}-A_{33}) & cA_{45}-sA_{34} & c^2A_{55}-2scA_{35}+c^2A_{33} & cA_{56}-sA_{36} \\ A_{16} & A_{26} & cA_{36}+sA_{56} & A_{46} & cA_{56}-sA_{36} & A_{66} \end{bmatrix}$$

- So
- $A'_{ij} = A'_{ji} = (c^2-s^2)A_{ij} + sc(A_{ii}-A_{jj})$ (1)
 - $A'_{ii} = c^2A_{ii} - 2scA_{ij} + c^2A_{jj}$
 - $A'_{jj} = s^2A_{ii} + 2scA_{ij} + c^2A_{jj}$
 - $A'_{ik} = A'_{ki} = cA_{ik} - sA_{jk} \quad \forall k \neq i, j$ (2)
 - $A'_{jk} = A'_{kj} = sA_{ik} + cA_{jk} \quad \forall k \neq i, j$ (3)
 - $A'_{kl} = A_{kl} \quad \forall k, l \neq i, j$

We can choose i and j so that A_{ij} is an entry off the leading diagonal with maximal modulus, and conjugate with a Givens matrix to make $A'_{ij} = 0$. From (1) above

$$\cos 2\theta A_{ij} + \frac{1}{2} \sin 2\theta (A_{ii} - A_{jj}) = 0$$

$$\Rightarrow \tan 2\theta = \frac{2A_{ij}}{A_{jj} - A_{ii}} \quad \text{with } \theta = 45^\circ \text{ (from the previous equation) if } A_{ii} = A_{jj}.$$

Clearly other entries have changed in moving from A to A' , and if the process is repeated it is unlikely that the newly-introduced zero will remain zero for long, so it is not obvious that we have done anything useful. We now show, however, that the sum of the squares of the entries off the leading diagonal has been reduced, and that the matrices generated by successive applications of the process are converging to a diagonal matrix. The diagonal entries of this must be the eigenvalues of A .

To show that the sum of the squares of the entries off the leading diagonal has decreased we can ignore all entries on one side of the leading diagonal because the matrix is symmetric, and clearly we can ignore the entries which do not change. We know that an entry with the greatest modulus has been changed to zero. This leaves $n-2$ pairs of entries, each one corresponding to a value of k , where $k \neq i, j$. From (2) and (3) above the sum of each of these pairs is

$$\begin{aligned} & (cA_{ik}-sA_{jk})^2 + (sA_{ik}+cA_{jk})^2 \\ &= c^2A_{ik}^2 - 2csA_{ik}A_{jk} + s^2A_{jk}^2 + s^2A_{ik}^2 + 2csA_{ik}A_{jk} + c^2A_{jk}^2 \\ &= A_{ik}^2 + A_{jk}^2 \end{aligned}$$

and so is unchanged. Thus the sum of all the squares of the entries off the leading diagonal has decreased.

2.2.3 Convergence

There are $\frac{n(n-1)}{2}$ entries on each side of the leading diagonal of which the one with the previously greatest modulus has been reduced to zero and the sum of the squares of the others is unchanged. So at each stage the sum of the squares is multiplied by a factor of at most $1 - \frac{2}{n(n-1)}$ and we have convergence at a useful rate.

A program was written to replicate the built-in function, though it was considerably slower because it was neither optimized nor compiled. This showed that it takes around 150 iterations to calculate the eigenvalues of a 10×10 adjacency matrix to 28 significant figures, which is PARI's default precision.

2.2.4 Eigenvectors

A huge advantage of Jacobi's algorithm is that the eigenvectors can be calculated at the same time as the eigenvalues with remarkably little extra effort, and this is always done by `qfjacobi`. This important feature is easily described in just seven lines, though it was not used at all in the production of this thesis.

The algorithm generates $G_n^T \dots G_1^T A G_1 \dots G_n \approx D$ where A is the adjacency matrix and the G_i are Givens matrices. D is the matrix with λ_i , the eigenvalues of A , on its leading diagonal and all other entries very close to zero. The eigenvectors of D are clearly the standard basis vectors \mathbf{e}_i (with component i equal to one and all other components equal to zero). That is $D\mathbf{e}_i = \lambda_i\mathbf{e}_i$.

So $G_n^T \dots G_1^T A G_1 \dots G_n \mathbf{e}_i = \lambda_i \mathbf{e}_i$ and since $(G_i^T)^{-1} = G_i$ we have $A(G_1 \dots G_n \mathbf{e}_i) = \lambda_i (G_1 \dots G_n \mathbf{e}_i)$.

So the eigenvectors of A are $G_1 \dots G_n \mathbf{e}_i$ which are just the columns of $G_1 \dots G_n$. This matrix is very easy to calculate while going through the algorithm to find D , without needing to store all the G_i .

2.2.5 Alternatives

The calculation of eigenvalues was so central to this thesis that two serious alternatives to `qfjacobi` were considered.

- Finding and solving the characteristic polynomial takes about 25% longer than `qfjacobi` for a typical 10×10 adjacency matrix (and doesn't find the eigenvectors).
- Also starting with the characteristic polynomial, Sturm's algorithm [A] is an extremely efficient way of finding how many eigenvalues are in a given half-open interval, even in cases where there are repeated roots, the corresponding factors of which have to be removed before the algorithm can be applied. But even when none of the eigenvalues or eigenvectors are required explicitly we frequently need the sum of the top and bottom eigenvalues as the easiest way of determining whether or not a graph is bipartite (in which case the sum is zero), otherwise some kind of colouring algorithm would have to be used, so `qfjacobi` is preferable.

2.3 PARI: advantages and problems

2.3.1 Advantages

The overwhelming advantage of PARI is that it works efficiently to any number of significant figures, limited only by memory and time constraints. It can calculate and display π to a million digits in 95 seconds. The other advantages are that it has a very large number of useful well-written built-in functions, and it's free.

2.3.2 File handling

There are four problems associated with PARI's handling of files.

First, files can be read only in their entirety — the whole thing at one go. The standard (and fairly obvious) business practice of reading a single record from a large sequential file, processing it, then replacing it by reading in the next record, is simply not possible. So it is necessary to keep track of the size of any output and to open a new output file, typically with essentially the same name but with an increased suffix number, whenever there is a danger of writing files that are too big to be read back in again at some future time.

Secondly it is not possible to delete or rename files from within PARI. So the standard procedure of

- inputting data
- processing it and outputting the results to a file with a temporary name
- deleting the original file and
- renaming the temporary file with the original file name

cannot be done. This, particularly when combined with the first problem, means that there is usually a very large build-up of intermediate files, but with a little care these can be confined to temporary directories and deleted *en masse* when a task is finished.

Thirdly PARI cannot create a file with nothing in it. So if, for example, output is directed according to some criterion to one of ten files named from **Output-01** to **Output-10**, any file to which no output was sent will not exist, rather than exist and be empty. Since PARI is also unable to check whether or not a file exists before attempting to open it, a program to read and process these files sequentially will crash if one is missing. This problem can be overcome by always writing something (the number zero was chosen for want of anything better) to every file which might exist, thus ensuring that it does exist. Programs which read files will then of course have to be written to ignore this first entry.

Finally, if PARI writes more than one object to a file, the file becomes a vector whose components are the objects which have been output to it. When the first (and perhaps the only) object is written to a file however, it is simply written as itself (meaning specifically that it is *not* the only component in a one-dimensional vector). When encountering a vector at the beginning of a file it is not always easy to tell whether the vector is the required object, or a simply a container whose first entry is the required object.

One solution to this problem is to write not just one zero to every file that might exist (as described in the previous paragraph), but two. The first zero will ensure that the file exists; the second will ensure that a container vector is created. Every file is now a vector. Vectors of length two represent what would normally be an empty file, and a non-empty file is a vector with more than two components, with the meaningful data beginning at the third.¹

2.3.3 Memory management

PARI can currently control just under one megabyte of memory, so using a computer with huge amounts of volatile memory does not enable the reading in of larger files.

More positively however, two copies of PARI (at least) can be run simultaneously with an increase in the time taken to perform calculations which is usually much less than 20%.

Many PARI functions can be applied to a much larger range of objects than might be expected. This appears at first to be very helpful, but it can actually lead to quite stunning inefficiency. Consider, for example, the **concat** function.

Most programming languages have some kind of concatenation function, usually limited to strings. So that, for example, a filename like “Output-” and a variable containing the number 10 can be concatenated to form the filename “Output-10” (usually with an intermediate function being necessary to turn the number 10 into the string “10”). With PARI the intermediate function is not necessary, and the **concat** function can be used with other objects. It is possible, for example, to **concat** an object to a vector, which increases the dimension of the vector by one and makes the object its last entry.

This sort of thing makes PARI amazingly flexible, but not necessarily usefully so. Consider the following example where we set up a 20×20 matrix with entries which are small integers:

```
m=matrix(20,20);for(i=1,20,for(j=1,20,m[i,j]=i-j))
```

Now we build a vector with 10,000 entries, each of which is a copy of this matrix. We do this in two ways. First we declare a vector of the required size and then fill in the entries:

```
v=vector(10000);for(i=1,10000,v[i]=m)
```

Secondly we start with a vector of length one, the only entry in which is a copy of the matrix, and **concat** the other entries, as one might wish to do when finding results one at a time:

```
w=vector(1);w[1]=m;m=w;for(i=1,9999,w=concat(w,m))
```

It should of course be obvious that computer programming is a combination of logic and engineering and is not at all surprising that the former is more efficient, but it is perhaps extremely surprising that the second method takes well over 7,000 times longer, comparable to the difference between a calculation taking one hour, or one year.

¹ Regrettably it took a very long time to evolve a mindset where this sort of surruptitious trickery is routinely used to fight the system. Perhaps too much study of mathematics discourages devious thinking.

2.4 An isomorphism algorithm

2.4.1 Introduction

The first big computational requirement in this thesis was for the calculation of eigenvalues, and an algorithm for doing this, which is very efficient and almost certainly the best available, has been described. The second big computational requirement was for the elimination of isomorphic graphs from sets of adjacency matrices. In this case, however, to quote [FSG], “As it could be expected, it does not exist an algorithm that is definitively better than all the others”. In view of this, and since the requirements here are slightly unusual (the elimination of isomorphisms from quite large numbers of quite small graphs) and since it promised to be an instructive and enjoyable exercise, it was decided that an algorithm should be specially written.

Given a complete set of pairs of associated vertices of two graphs then the two graphs are isomorphic if and only if there is an edge between two vertices of one when and only when there is an edge between the associated vertices of the other. The obvious algorithm to find whether one graph is isomorphic to another is to fix the labelling of the vertices of one graph and permute the labelling of the vertices of the other in every possible way, checking for isomorphism each time. Clearly one can go on to the next permutation as soon as any contradiction is found, and finish as soon as any isomorphism is found. Even so this is almost never efficient and very rapidly becomes computationally infeasible.

2.4.2 The Distance Matrix and the Correspondence Vector

The Distance Matrix is a table of information about a graph which is stored as a matrix so that use can be made of PARI’s many simple but efficiently-written functions for manipulating matrices.

Row r of the Distance Matrix initially corresponds to the same vertex of the graph as row r of the adjacency matrix. The entry in column c of a row in the Distance Matrix is the number of vertices of the graph which are distance c from the vertex which the row represents.

The rows of the Distance Matrix are now sorted in some way which we choose to regard as standard.

Clearly this breaks the connection between the rows of the Distance Matrix and the rows of the adjacency matrix, so before sorting a Correspondence Vector is created, the entries of which initially run simply from 1 to n , where n is the number of rows of the Distance and adjacency matrices. When the rows of the Distance Matrix are sorted, the entries of the Correspondence Vector are permuted in exactly the same way, so that the correspondence between the rows of the Distance Matrix and the rows of the adjacency matrix is not lost.

This is done for the following reasons:

- The Distance Matrix contains a lot of information about its graph which can be calculated once only, stored efficiently with the adjacency matrix, and used in thousands of tests for isomorphism without the need for recalculation.
- Equality of sorted Distance Matrices is clearly a necessary condition for isomorphism, so a huge number of non-isomorphisms can be established by a simple comparison with no further calculation.
- The Distance Matrix partitions the rows of a graph into equivalence classes which have the same entries in the same order. Any corresponding vertices in two isomorphic graphs must belong to the same equivalence class, so the more equivalence classes there are, the more the number of permutations needed to find or disprove an isomorphism is (massively) reduced.

If we look at the top Distance Matrix and Correspondence Vector in the example on the next page, we see that the first graph has 1 vertex of valency 6, and the 5 vertices which are not adjacent to it are distance 2 away. The vector tells us that this corresponds to the 6th row/column of the adjacency matrix. The two adjacency matrices shown have the same Distance Matrix, which is clearly a necessary condition for isomorphism of their graphs, which are in fact isomorphic in this case.

Rows 1, 6, 10, 11 and 12 of the Distance matrices are unique within each matrix, which means, using the Correspondence Vectors, that if there is an isomorphism then vertex 6 of the first graph must correspond to vertex 9 of the second graph, vertex 3 of the first graph must correspond to vertex 8 of the second graph and so on. We call these correspondences unavoidable. Similarly the set of four vertices $\{4, 9, 10, 12\}$ of the first graph corresponds to the set of four vertices $\{4, 5, 11, 12\}$ of the second graph (in some order) because they all have the same entries in the Distance Matrices, and the set of three vertices $\{1, 7, 11\}$ in the first graph must correspond to the set of three vertices $\{1, 2, 6\}$ in the second. So in this particular case we now know that there are only $4! \times 3! = 144$ possible pairings of vertices for isomorphism and a full edge test in each case is certainly computationally feasible.

This partitioning of vertices doesn't always happen; there are 34 cyclotomic/Salem graphs with number of vertices between 2 and 10 which have every row of their Distance matrices the same. They include the cyclic graphs, the complete graphs and the Petersen graph.

Example of the use of the algorithm to establish isomorphism



Adjacency Matrices											Distance Matrices			Correspondence Vectors	
0	0	1	0	0	0	1	0	1	1	0	0	6	5	0	6
0	0	0	1	0	0	0	0	1	1	0	1	5	6	0	4
1	0	0	1	1	1	0	0	0	0	0	0	5	6	0	9
0	1	1	0	0	1	0	0	0	1	1	0	5	6	0	10
0	0	1	0	0	0	1	0	0	0	1	0	5	6	0	12
0	0	1	1	0	0	0	1	1	1	0	1	4	7	0	3
1	0	0	0	1	0	0	0	0	1	0	1	4	6	1	1
0	0	0	0	0	1	0	0	0	0	0	0	4	6	1	7
1	1	0	0	0	1	0	0	0	0	1	1	4	6	1	11
1	1	0	1	0	1	1	0	0	0	0	0	4	5	2	2
0	0	0	1	1	0	0	0	1	0	0	1	3	6	2	5
0	1	0	0	0	1	1	0	1	0	1	0	1	5	5	8
0	0	0	1	0	0	1	0	0	0	1	1	6	5	0	9
0	0	0	1	1	1	0	1	0	0	0	0	5	6	0	4
0	0	0	1	1	0	0	0	0	0	1	1	5	6	0	5
1	1	1	0	0	0	0	0	1	0	1	0	5	6	0	11
0	1	1	0	0	1	0	0	1	0	0	1	5	6	0	12
0	1	0	0	1	0	1	0	0	0	1	0	4	7	0	8
1	0	0	0	0	1	0	1	0	0	0	0	4	6	1	1
0	1	0	0	0	0	1	0	1	0	0	1	4	6	1	2
0	0	0	1	1	0	0	1	0	1	1	1	4	6	1	6
0	0	0	0	0	0	0	0	1	0	0	0	4	5	2	3
1	0	1	1	0	1	0	0	1	0	0	0	3	6	2	7
1	0	1	0	1	0	0	1	1	0	0	0	1	5	5	10

Assignment Table

Matrix 1	1	2	3	4	5	6	7	8	9	10	11	12
Matrix 2		3	8		7	9		10				
Level	X	0	0	X	0	0	X	0	X	X	X	X

Permutation Table

Level	1	2	3	4	5	6	7
Maximum	3	2	1	0	2	1	0
Counter	0	0	0	0	0	0	0
Graph 1 vertices (fixed)	4	9	10	12	1	7	11
Graph 2 vertices (cycle)	4	5	11	12	1	2	6



2.4.3 The Assignment table and the Permutation Table

Now we set up an Assignment Table and a Permutation Table.

The number of columns in the Assignment Table (apart from the labels in the first column of the printed version of course) is equal to the number of vertices in the graphs being compared. We assume that it has already been checked that this is the same for both.

- The first row consists of the numbers of the vertices of the first graph (the rows/columns of the first matrix), in order. These are simply the numbers of the columns of the table so they are not actually stored in computer memory — but it is clearly convenient to be able to see them in the printed version.
- The second row consists of the vertices of the second graph which correspond to the vertices of the first. Initially, as in the example, only the unavoidable assignments are filled in. If an isomorphism is found then, when the algorithm terminates, this row gives the precise correspondences which constitute the isomorphism. At all other times all the assignments, apart from those which are unavoidable, are provisional; they might be undone at any time.
- The third row facilitates this potential undoing. It shows the order in which assignments have been made so that they can be undone in reverse order if necessary. A zero means that an assignment is unavoidable and cannot be undone. An X means that either no assignment has been made or that any assignment made has been undone, so there is no current valid assignment. Computer memory can never be blank of course; an X means that whatever is in the memory corresponding to the cell above the X is not of any importance. As non-unavoidable assignments are made they are numbered, starting from 1, in this row.

The number of columns in the Permutation Table (apart from the labels in the first column of course) is equal to the number of vertices for which the assignments are not unavoidable. They are grouped into blocks separated by heavy vertical lines, and the sizes of the blocks depend on the numbers of identical rows in the Distance Matrices which is the size of the equivalence classes which have more than one member. In the example there are 4 rows with 5, 6, 0 in the Distance Matrix, and 3 rows with 4, 6, 1, so the blocks have widths of 4 columns and 3 columns.

The elements in the bottom row of the table are to be permuted cyclically within their block. A pointer which points initially to the first column is shown below the table in the example. The cycles extend from the column to which the pointer is pointing to the end of the block in which the pointer is currently situated. The cycles in the example are shown at the bottom of the page. The first cycle would currently consist of the numbers 4, 5, 11 and 12. The second cycle, as the pointer moves to the right, would currently consist of the numbers 5, 11 and 12, and the third would currently consist of the numbers 11 and 12. Another cycle would permute a single number, so it would do nothing and is not shown.

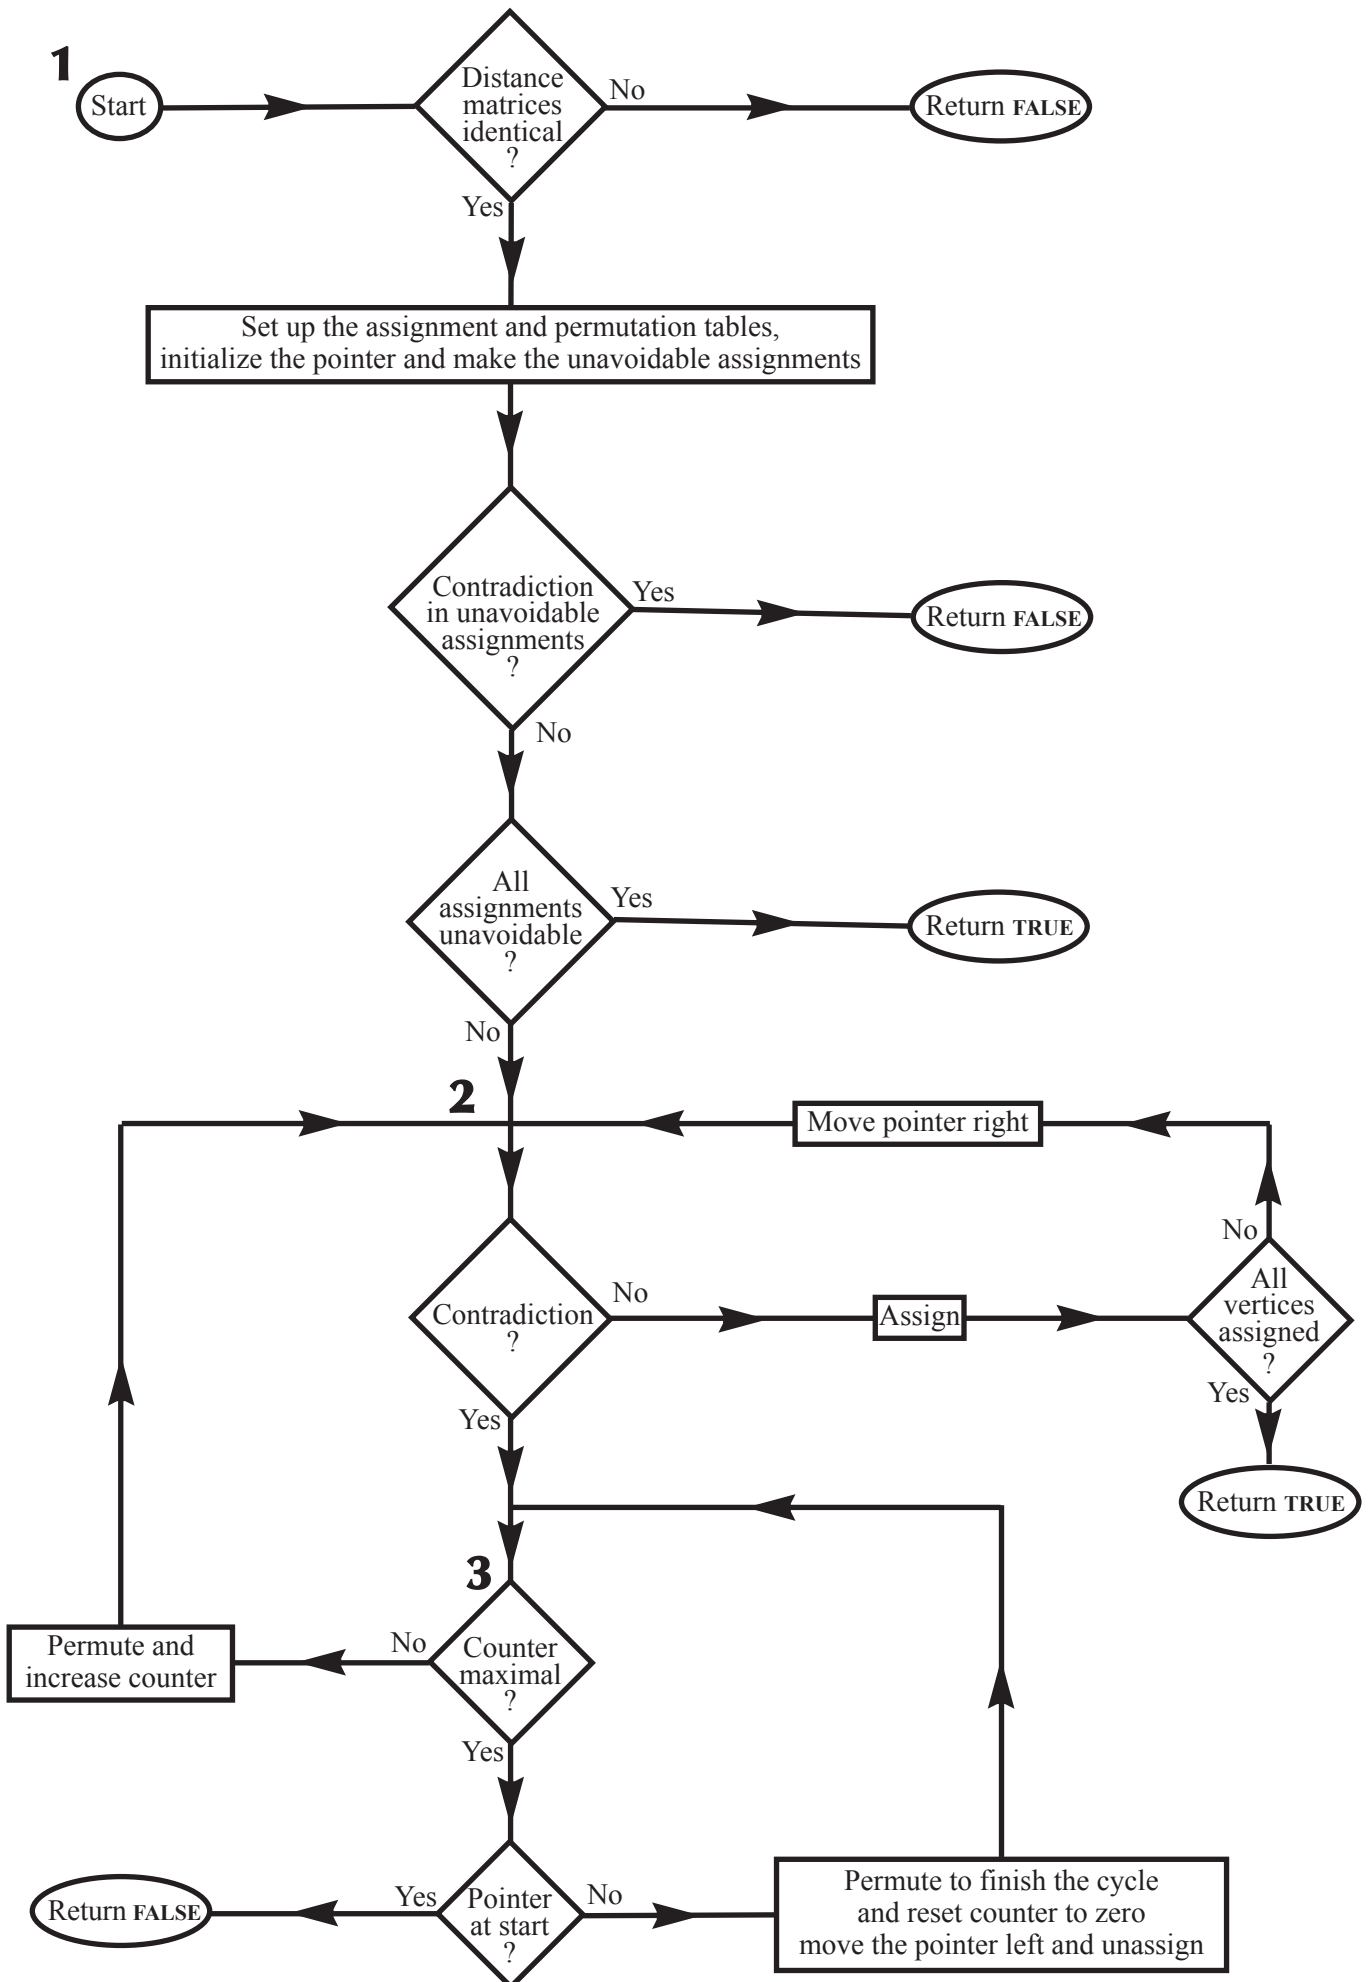
Then the pointer moves to the next block where there are two cycles currently consisting of the numbers 1, 2, 6 and 2, 6.

The Permutation Table is made up of the following rows:

- The first row is the order in which assignments will be made, used to fill in the third row of the Assignment Table. These are simply the numbers of the columns of the table so they are not actually stored in computer memory — but it is clearly convenient to be able to see them in the printed version.
- The second row shows the maximum number of cyclic permutations which are possible of the numbers in the columns from the current column to the column at the end of the current block before repetitions start. These numbers simply count down to zero within each block. The cells containing zero ensure that a cycle consisting of a single number is never permuted.
- The third row is initially set to zero and shows how many of the above cyclic permutations have been carried out. If this number were allowed to exceed its maximum in the cell above we would be carrying out pointless repetitions.
- The fourth row consists of vertices in the first graph which do not have a unique row in the Distance matrix. Within each block the row of the Distance Matrix corresponding to each vertex is the same. Within each block the vertices could be in any order, but for computational convenience they will always be in numerical order. The row never changes.
- The last row consists of vertices in the second graph which do not have a unique row in the Distance Matrix. Within each block the rows of the Distance Matrix corresponding to each vertex are the same, and also the same as the rows of the Distance Matrix corresponding to the vertices of the first graph in the cells above. Within each block the vertices could start in any order, but for computational convenience they will always start in numerical order. This is the row which is subject to the cyclic permutations described above.

An assignment is made by taking the column of the Assignment Table corresponding to the Graph 1 vertex above the pointer, placing the Graph 2 vertex above the pointer in the second row of the Assignment Table, and placing the Level above the pointer in the third row (the Level row) of the Assignment Table. An unassignment is made by taking the column of the Assignment Table corresponding to the Graph 1 vertex above the pointer, and placing an X in the third row (the Level row) of the Assignment Table. This means that the number in the second row is not significant — it is not necessary to delete it.

Flow chart for the isomorphism algorithm



2.4.4 Contradictions

The crucial concept of a contradiction must now be made precise.

When making a new (provisional) assignment we use the assignment table to take in turn each pair of vertices which have already been assigned (both unavoidably and provisionally).

- We use the Correspondence Vector and the Adjacency matrix of the first graph to find whether or not the already-assigned vertex of the first graph is adjacent to the vertex of the first graph in the new provisional assignment.
- We use the Correspondence Vector and the Adjacency matrix of the second graph to find whether or not the already-assigned vertex of the second graph is adjacent to the vertex of the second graph in the new provisional assignment.

The adjacency (or not) must be the same in each case or there is a contradiction because the new provisional assignment cannot logically be made.

The flow chart of the algorithm can now be described.

- 1 If the Distance Matrices are not identical return **FALSE**. Set up the Assignment and Permutation tables, initialize the pointer, and make the unavoidable assignments. If there is a contradiction in the unavoidable assignments return **FALSE**. If all the assignments are unavoidable and there is no contradiction return **TRUE**.
- 2 If there is a contradiction between the vertices of Graph 1 and Graph 2 which are above the pointer **GOTO** 3. If there is not a contradiction between the vertices of Graph 1 and Graph 2 which are above the pointer then make the provisional assignment and if all the assignments have been made return **TRUE**. Otherwise move the pointer to the right and **GOTO** 2.
- 3 If the counter above the pointer is not maximal then carry out a permutation, increase the counter and goto 2. If the counter above the pointer is maximal and the pointer is back at the start then we have exhausted every possibility so return **FALSE**. Otherwise we have only exhausted the current cycle so carry out a final permutation to get the cycle back to its beginning and reset the counter above the pointer to zero. Move the pointer to the left to return to the previous cycle and unassign the provisional assignment which is now above the pointer because it led to an unavoidable contradiction (an infeasible cycle). **GOTO** 3.

Some Salem graphs with unusual distance tables I




Figure 1
The smallest pair of non-isomorphic Salem graphs with identical distance tables
which have no repeated rows, so all assignments are unavoidable




Figure 2
The smallest pair of non-isomorphic Salem graphs with identical distance tables
in which every row is the same; namely 3 2 0 0 0


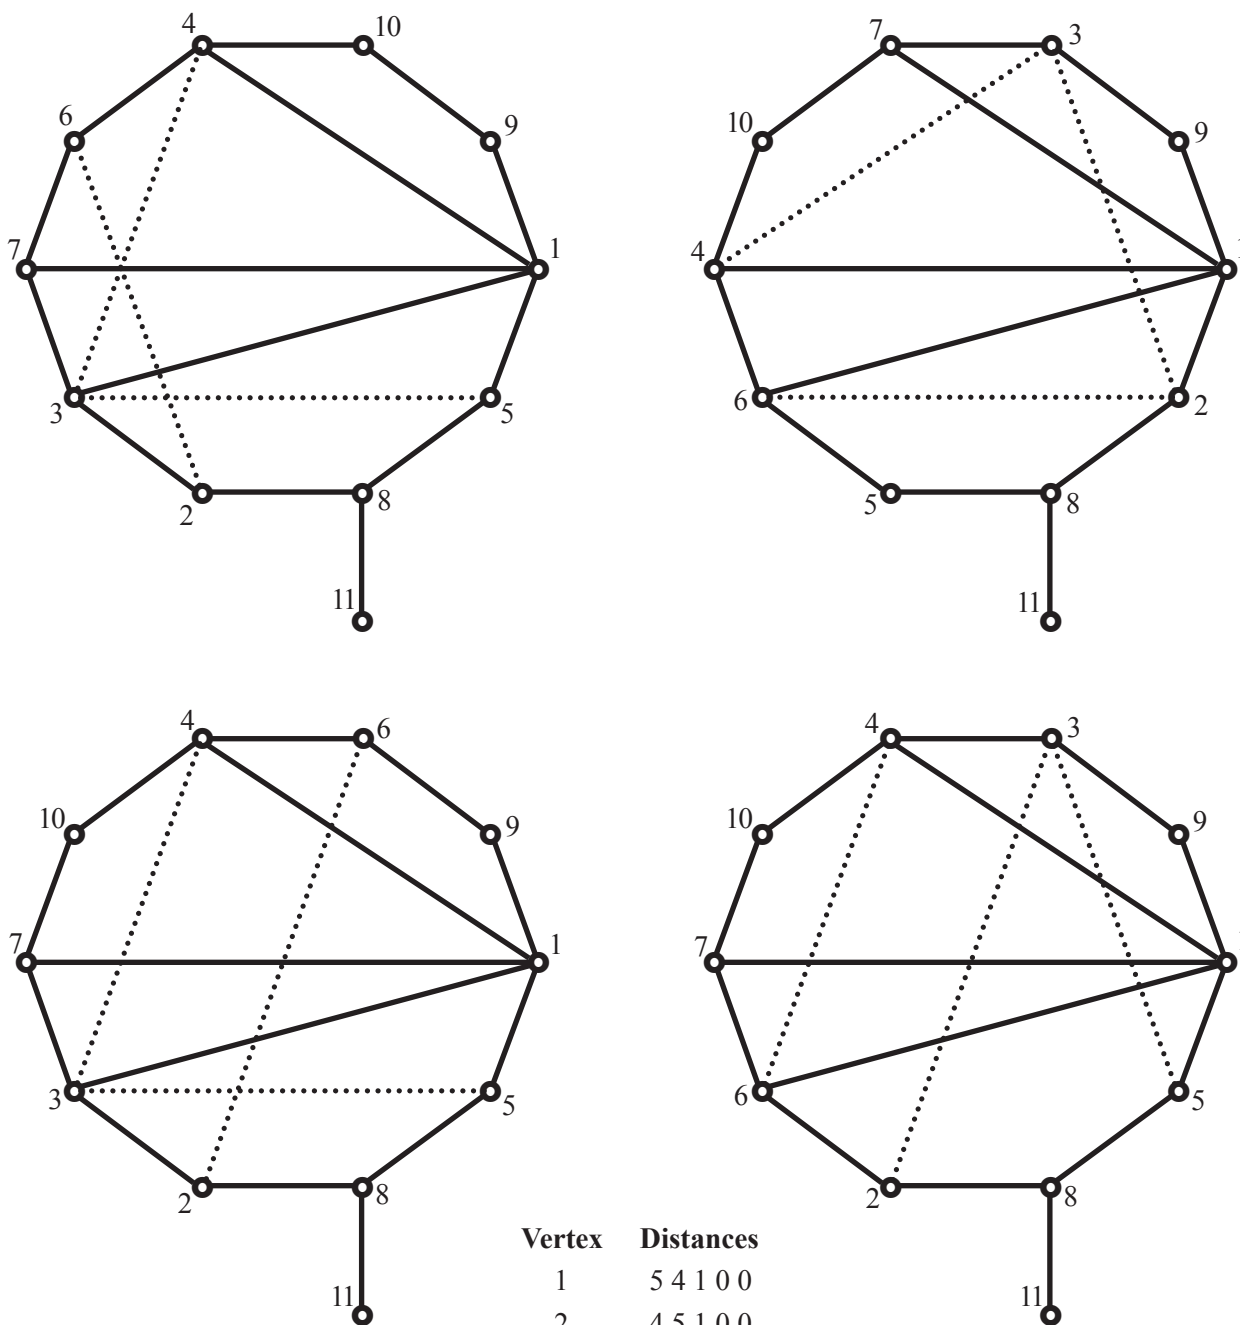


Figure 3
The smallest set of three non-isomorphic Salem graphs with identical distance tables
in which every row is the same; namely 4 3 0 0 0 0 0

Some Salem graphs with unusual distance tables II



Vertex	Distances
1	5 4 1 0 0
2	4 5 1 0 0
3	4 4 2 0 0
4	4 4 1 1 0
5	3 5 2 0 0
6	3 4 3 0 0
7	3 4 2 1 0
8	3 3 3 1 0
9	2 4 3 1 0
10	2 3 3 1 1
11	1 2 3 3 1

Figure 4
A set of four non-isomorphic Salem graphs with identical distance tables
which have no repeated rows, so all assignments are unavoidable

2.5 Some Salem graphs with unusual distance tables

2.5.1 Introduction

In this section we consider Salem graphs with identical distance tables at the two extremes:

- those with distance tables in which every row is different, which the algorithm always finds extremely straightforward, and
- those with distance tables in which every row is the same. When the graphs are isomorphic there is usually so much symmetry that the algorithm finds one of the automorphisms quickly. When the graphs are not isomorphic the distance tables are of no use and many permutations of vertices have to be tried.

Some of the edges in the examples are drawn as dotted lines. This is just an attempt to present the graphs in a more memorable way; the dotted edges have exactly the same status as the solid edges.

2.5.2 Every row different

If two graphs have identical distance tables in which every row is different then all the assignments are unavoidable. If the vertices of one graph are numbered in some way then there is no choice in how to number the vertices of the other graph. The concept of the distance table is at its most useful in this case because no permutations are required. To establish if two such graphs are isomorphic it is simply necessary to check that there is an edge between two vertices of one graph if and only if there is an edge between the same two vertices in the other graph.

The smallest of these pairs of graphs, perhaps surprisingly the only one with fewer than ten vertices, is shown in figure 1. The algorithm can tell them apart because there is, for example, an edge between vertices 2 and 4 in one graph but not in the other. It is often interesting to imagine a pair of non-isomorphic graphs being presented as a collection of rings tied together with bits of string, and to give a simple rule for a human to distinguish between them without needing to use a distance table or number the vertices. In this case we simply locate the three vertices with valency two. In one graph one of these vertices is adjacent to both the others, in the other graph this is clearly not the case.

There are five such pairs of graphs with ten vertices (with 14, 18, 20, 20 and 22 edges) and eighteen such pairs with eleven vertices. In addition to these eighteen pairs the smallest such set of more than two graphs is a remarkable set of four shown in figure 4. They can be thought of as an underlying graph with 11 vertices and 14 edges, with just three extra edges, shown dotted, added differently in each case. The top left graph is easily distinguished because the two vertices with valency two are adjacent. To distinguish the others consider the four vertices with valency three: in one graph three of the four are in a row, in another two of them are separated from the others, and in the last graph they are in two adjacent pairs.

2.5.3 Every row the same

Now we turn to sets of non-isomorphic Salem graphs with identical distance tables in which every row is the same. Clearly in these cases the distance table is completely useless in trying to find isomorphisms. However, given two isomorphic copies of one of these graphs, such as the Petersen graph, the order of the automorphism group is so large that it is usually possible to find an isomorphism mapping without going through an enormous number of vertex permutations. The slowing down of the algorithm seems to occur mainly with graphs which are “almost isomorphic”, meaning that they are not isomorphic but have one or more large sets of similar vertices in the permutation table, an example of which is described in the next section.

The only pair of such graphs with fewer than eight vertices is shown in figure 2. These are $K_{3,3}$ and the triangular prism. They are very easily distinguished by eye because the prism is planar with every vertex being a vertex of a triangle, while $K_{3,3}$ is bipartite (with no triangles) and is famously not planar.

2.6 The structure of two “almost isomorphic” graphs

2.6.1 Two graphs in E_8

Of all the graphs to which the algorithm has been applied which have arisen naturally (so excluding those described below which are constructed with the sole purpose of being difficult for the algorithm to distinguish) two stand out as requiring an unusually long time (well over a minute) for their lack of isomorphism to be established. We look at these in detail for several reasons:

First to see why the algorithm is inefficient in this case, how to construct an even “worse” case, and how the algorithm might be improved (or at least changed) to deal with this case. Secondly because the structure relates back in a very pleasing way to the small graphs with unusual distance tables described in the previous section, and finally because the two objects and the relationship between them are examples of great mathematical beauty.

The two graphs are both maximal subgraphs of E_8 with 30 vertices. The following is a description of both graphs until otherwise stated.

Imagine 12 of the vertices in a horizontal regular dodecagon with one vertex above its centre and one below. The other 16 vertices can be thought of as “evenly distributed” on a sphere underneath. The top vertex is connected to every vertex of the dodecagon. The dodecagon consists of the vertices of the cocktail party graph CP_{12} , “almost” a copy of K_{12} , with just six independent edges missing, which can be thought of as the six between opposite vertices, so the complement consists of six copies of K_2 . The vertex below the dodecagon is connected to every vertex of the dodecagon and to every vertex of the sphere. Each vertex of the dodecagon is connected to eight vertices on the sphere, and each vertex

The structure of two “almost isomorphic” maximal subgraphs in E_8

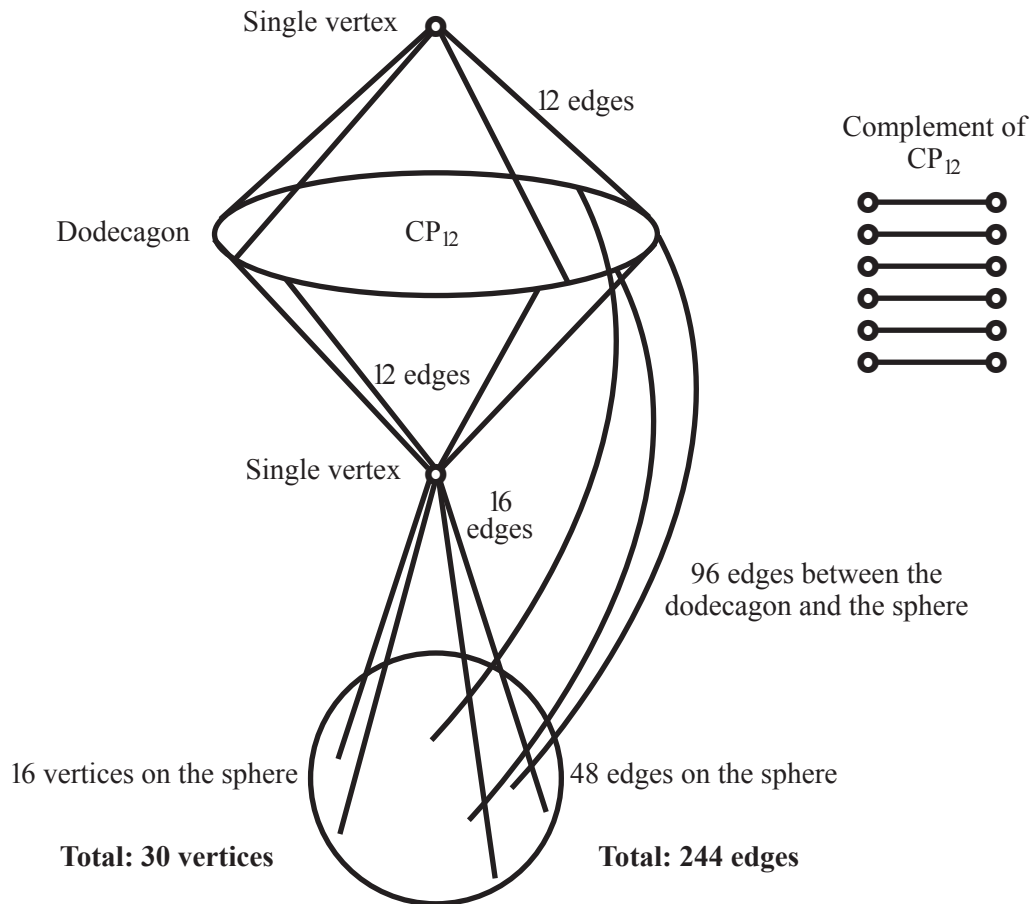


Figure 1




Figure 2




Figure 3




Figure 4




Figure 5




Figure 7




Figure 6

on the sphere is connected to six of the others. All this is shown in figure 1. The differences between the two graphs lie solely in the detail of the pattern of edges on the sphere.

Figure 2 shows the pattern of edges on the sphere of the first graph. Following any four edges in the same direction gets you back to where you started. Now consider the vertices inside the grey oval (the one which is not dashed). These eight vertices and the edges between them can be drawn as in figure 3, which is familiar from the previous section. The other eight vertices which are not in the grey (not dashed) oval and the edges between them can also be drawn in exactly the same way. All this is shown in figure 4, which also shows four of the 16 edges between the two cubes. These 16 edges can be described as follows:

For each vertex of the first (upper left) cube, identify the vertex which can be reached along a dotted edge, and the vertex which is diagonally opposite through the cube. The vertex has edges going to the two corresponding edges in the second cube. For example start with vertex 3. The dotted edge from 3 goes to 14 and the vertex diagonally opposite through the cube is 10. 14 and 10 correspond to 8 and 4 in the second cube, so 3 is connected to 4 and 8. Similarly 7 is connected to 12 and 8. The other 12 edges between the two cubes are not shown.

So figure 4 (along with the those other 12 edges) contains exactly the same information as figure 2, describing the vertices and edges on the sphere, but presented in a way which at first appears unnecessarily much more complicated. Now we consider the edges between the sphere and the rest of the structure. We already know that all 16 vertices are connected to the vertex below the dodecagon, but the edges between the dodecagon and the sphere are much more interesting. The eight edges from one vertex of the dodecagon to the sphere go to the eight vertices of one cube, and the eight edges from the opposite edge in the dodecagon (the only vertex to which the previously-considered vertex is not connected within the dodecagon) go to the eight vertices of the other cube.

Clearly there are six different pairs of opposite vertices within the dodecagon, and indeed there are six possible ways of choosing the vertices of the sphere to form a diagram similar to figure 4. Returning to figure 2, we have considered the vertices inside the unbroken grey oval, and the rest. We could also take the vertices inside the dashed grey oval, and the rest. The major axes of both these ovals are clearly parallel to one of the sides of each of the equilateral triangles which make up figure 2, so clearly it would be easy to draw similar pairs of ovals with major axes parallel to each of the other two sides of the equilateral triangles, giving the six possible constructions of figure 4.

This completes the description of the first graph.

The essential difference between the two graphs is that figure 6 (also familiar from the previous section) replaces figure 3. There is just one detail of interest. In figure 5 (which corresponds to figure 2 for the previous graph) the two grey ovals can only be rotated in only one (obvious) way to provide two more of the required six different selections of eight vertices for the cube, rather than four. Figure 7

shows the final two selections of the required six. It is easy to see (though much less easy to spot in the first place) that all six are equivalent.

Now we can see why the isomorphism algorithm is not efficient in this case.

The two vertices above and below the dodecagon are unavoidable assignments so take up almost no time when the algorithm is executed. The 12 vertices of the dodecagon have 20 neighbours, so come before the 16 vertices of the sphere (which have only 13 neighbours) in the permutation table. The cocktail party graph has so much symmetry that many permutations do not throw up a contradiction. The unavoidable contradictions are thrown up by the sphere, and one of these have to be found for each acceptable permutation of the dodecagon. This is why it takes over a minute to establish that the two graphs are not isomorphic.

2.6.2 A worse case

Each vertex of the dodecagon has 20 neighbours. To construct a “worse” example, all we need to do is to add a highly symmetrical structure which will appear earlier in the permutation table (meaning that it must have more than 20 neighbours) without messing up the symmetries of the previous graph. So we take a copy of K_{21} and place it at the top of the previous structures with each vertex of the K_{21} connected to the vertex which was previously at the top. For each of the $21!$ symmetries of the K_{21} a contradiction has to be found in what was the previous graph. Since it took over a minute to distinguish between the previous graphs it will now take over $21!$ minutes; greater than the age of the universe, to distinguish between the new graphs. The number of steps is finite, however, so it’s probably just about allowable to continue to call the method an algorithm.

In practice the algorithm has worked always adequately and almost always very well, so no attempt has been made to enable it to manage artificially difficult cases. It might be possible to put the blocks of similar vertices in the permutation table in a different order so that contradictions are found earlier, but if you understand the structure of the graphs well enough to decide which order of the blocks is best, you probably don’t need an isomorphism algorithm anyway. Probably it would be best to collect more information about the graphs which is invariant under isomorphism so that the blocks in the permutation table are split into smaller components.

Chapter 3

A catalogue of small Salem graphs

Without doubt many of the small graphs considered here must have been looked at before, in particular there have been a tiny number of classifications of some very special families (see [GM], to which the present author made a few small contributions, and [MS]). This, however, is the first systematic study of Salem graphs starting with the smallest and building up. In particular, the numerical results of how many Salem graphs exist with a given number of vertices listed in §3.2 were not previously known.

3.1 Building the catalogue

3.1.1 Introduction

When starting to investigate Salem graphs the most obvious question by far is how many there are with some given number of vertices. This chapter describes how that question has been answered, at least for rather small numbers, and how a catalogue of those graphs has been built.

3.1.2 Interlacing

Cauchy's interlacing theorem is the key to finding Salem graphs vastly more efficiently than by a brute-force search of all graphs.

Adjacency matrices are special cases of Hermitian matrices, which have real eigenvalues.

Cauchy's Interlace Theorem:

Let H be a Hermitian matrix with eigenvalues $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ and

H' a principal submatrix with eigenvalues $\mu_1 \leq \mu_2 \leq \dots \leq \mu_{n-1}$.

Then $\lambda_1 \leq \mu_1 \leq \lambda_2 \leq \mu_2 \leq \dots \leq \mu_{n-1} \leq \lambda_n$.

The proof given in [H] uses only undergraduate mathematics and is based on the intermediate value theorem. Surprisingly it has only been available since 2004. Previous proofs, not difficult to find, depended on other theorems which are themselves very far from obvious. The present author has a strong preference for the former, particularly since it occupies little more than a single page.

Cauchy's theorem is invaluable for this reason:

All connected graphs can be "grown" from the graph which consists of a single vertex by adding vertices one at a time, together with all the edges which go between the new vertex and the already-existing vertices. The crucial point is that all edges between previously-existing vertices are already present and never added later. So at each stage the old adjacency matrix is a principal submatrix of the new adjacency matrix, and Cauchy's interlace theorem applies. This means that once a grown graph has eigenvalues which prevent it from being either cyclotomic or Salem, no matrix grown from it can ever return to being either cyclotomic or Salem.

The strategy, then, is as follows:

Enter the two connected cyclotomic and Salem graphs with three vertices “by hand”. (Starting with one or two vertices saves no significant effort and makes the coding awkward.) Then at each stage grow all the connected graphs from the previous stage in every possible way, save those which are cyclotomic or Salem and remove isomorphic repeats.

The number of connected graphs with n vertices grows so rapidly with n that many exhaustive investigations of graphs rapidly become computationally infeasible, but if we restrict our attention to cyclotomic and Salem graphs, it is reasonable to hope (and it is in fact the case) that the terminal leaves in the growth tree (which appear when a grown graph is neither cyclotomic nor Salem) will restrict the number of graphs to such an extent that some exhaustive investigations can go considerably further than would otherwise be possible.

3.1.3 Standard graph format

Terabyte hard drives have recently become affordable, so drive space is, in crude terms, a less valuable resource than time. Storing graphs in a way which is relatively heavy on memory but speeds up the elimination of isomorphisms would seem to be sensible.

It was decided to store graphs as vectors with five components as follows:

- First the adjacency matrix.
- Next the Distance Matrix and Correspondence Vector as described in section 2.4.2.
- Finally it was felt that it would be advantageous to record whether or not the graph is bipartite, and whether it is cyclotomic or Salem.

3.1.4 Standard file format

The limitation on file size imposed by the fact that files can only be read in their entirety means that there must be a carefully-managed limit on the number of graphs which a file contains. A convenient number when dealing with the larger of the small Salem graphs being considered here was found to be 500.

To deal with the other problems of PARI described in section 2.3.2 the standard file format was chosen as follows:

Every file consists of a vector with at least 2 and not more than 502 components. The first two components are dummy entries (in fact the number zero). All other entries are graphs represented in the standard graph format described above. Files frequently exist in groups which are named with a text prefix followed by a number from 1 up to some limit. All files with numbers between 1 and that limit exist, even if they contain no “real” data.

3.1.5 Growing graphs

At each stage every graph from the previous stage is grown as follows:

Add a row and column to the adjacency matrix representing the new vertex. Fill in the row and column (which clearly must be the same in a symmetric matrix) in every possible way, one after another. In each case, if the resulting matrix is cyclotomic or Salem then expand it to standard graph format as described above and save it.

The possible ways of filling in the row and column are simply the binary numbers of the appropriate length, excluding zero which would create a disconnected graph, and excluding the odd numbers (which end in a 1 in binary) which would put a 1 on the leading diagonal and indicate a loop which is clearly not allowed in a “proper” graph.

In every case the eigenvalues are found by the `qfjacobi` function, which is absolutely crucial to the whole process. It is the only PARI function used with underlying details of operation which are not reasonably obvious, and it has been described in detail in section 2.2. The eigenvalues establish whether the graph is cyclotomic or Salem, and whether it is bipartite, since in this case the top and bottom eigenvalues add to zero because in a bipartite graph the eigenvalues are symmetric about zero [Bi p53]. For computational purposes we look for top and bottom eigenvalues adding to epsilon, where epsilon is “very close” to zero. 10^{-20} was chosen.

The details of the rest of the code are not of sufficient interest to be worth describing here.

3.1.6 Removing isomorphisms

With almost 100,000 graphs in files of 500, when only one of those files can be loaded at a time, the problem of comparing all possible pairs of graphs to eliminate isomorphisms is obvious.

It was found highly effective to do a preliminary separation of the graphs into relatively small collections with the same Distance Table, since all isomorphisms would then reside within these collections and not between them. The isomorphism algorithm could then operate easily on each of those collections which consisted of a single file.

In the relatively few cases in which a collection required more than one file the isomorphisms were eliminated within each file, the files were then merged, and the process repeated. This was tedious but not difficult. Clearly the problem would become much harder if the number of graphs were to approach a million, but equally clearly this situation will almost certainly appear amusingly dated within a few years.

3.1.7 Pausing

A final small but important aspect of the programming is the potential to pause. It is difficult to dedicate a much-used computer to a single task for well over a week, and running two programs simultaneously is very undesirable since at this time computer multi-tasking is still in its infancy; a crash in one task

will almost always freeze the whole computer and terminate all the other tasks, and real-time operations such as burning a DVD or playing a video clip do not yet share time happily. There are plenty of convenient points both in the growing of graphs and in the removal of isomorphisms where the state of the task in hand can easily be saved and the program temporarily halted. The obvious thing to do is to check at each convenient point whether some key has been pressed, but there is no way of doing this from within PARI. There is, however, a timer, and it is easy to check at each convenient point whether or not a time specified at the start of the program has been exceeded, and proceed accordingly. This is far from ideal, but at least enables good use to be made of some substantial chunks of time when the computer is not otherwise required, for example at night.

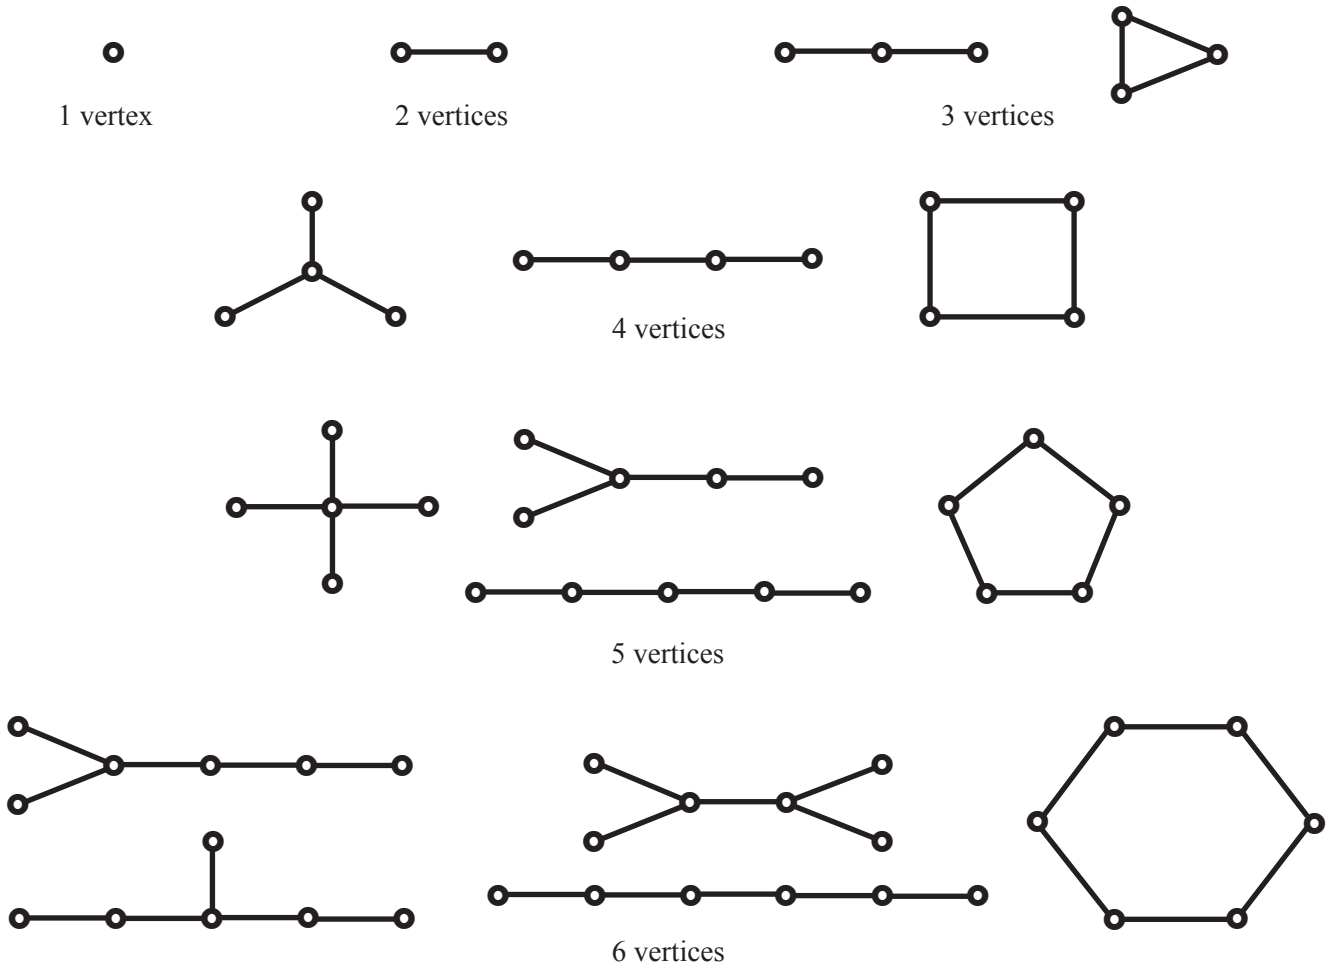
3.2 Numerical results

The counts which have been established so far are an excellent illustration of the often-encountered so-called “combinatorial explosion”. They are as follows:


Table 3.1 Numbers of cyclotomic and Salem graphs with a given number of vertices

Vertices	Cyclotomic Graphs	Salem Graphs
1	1	0
2	1	0
3	2	0
4	3	3
5	4	16
6	5	74
7	6	324
8	6	1496
9	5	6201
10	4	20723
11	4	58952

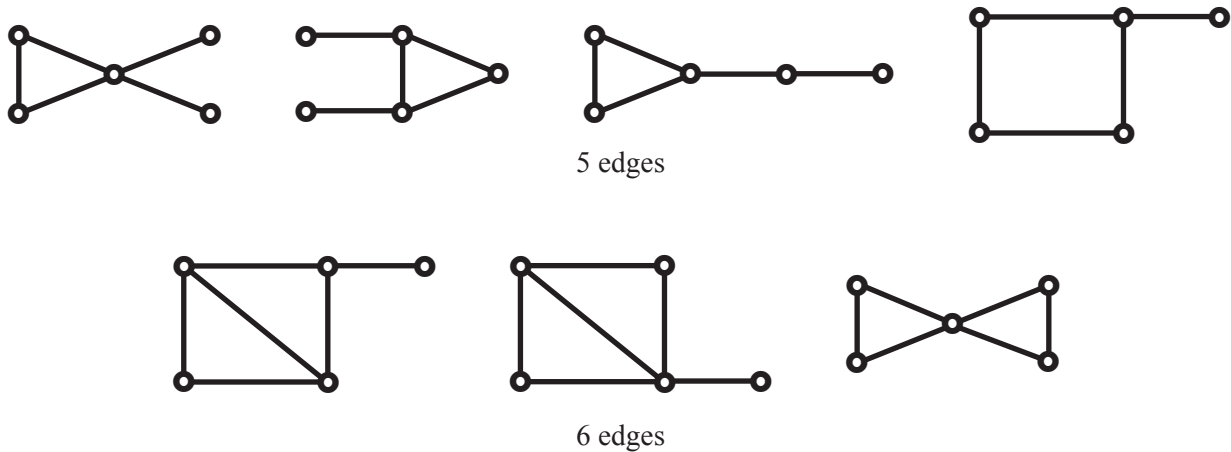
3.3 The 16 cyclotomic graphs with 6 vertices or fewer





3.4 The 3 Salem graphs with 4 vertices






3.5 The 16 Salem graphs with 5 vertices






3	1
3	1
2	2
2	2
2	2



6 edges with the same distance table




7 edges



8 edges





9 edges










10 edges



3.6 The 74 Salem graphs with 6 vertices




5 edges




6 edges






7 edges




7 edges with the same distance table

3	2	
3	2	
3	2	
2	2	1
2	2	1
1	2	2




8 edges




4	1
4	1
2	3
2	3
2	3
2	3

and




3	2	
3	2	
3	2	
3	2	
2	2	1
2	2	1


8 edges with the same distance tables





9 edges




9 edges and the same distance tables




10 edges




11 edges




12 edges



13 edges



14 edges



15 edges

Chapter 4

GRINs

Apart from the listing of the 18 graphs below, everything in this chapter is original.

4.1 The 18 GRINs


4.1.1 Definitions

The *index* of a graph is the greatest eigenvalue of its adjacency matrix. A graph is said to be *minimal* with respect to some property if none of its induced subgraphs have that property. There are precisely 18 graphs which are minimal with respect to the property of having index greater than two. These will be called *GRINs* (**G**raphs with **I**ndex **G**reater than two which are **m**INimal).


GRINs are important here because, by interlacing, every Salem graph has at least one GRIN as an induced subgraph.

The graphs have been taken from [CvRo], but have been renamed because of the importance here of whether or not they are bipartite.


4.1.2 Figure 1: list of GRINs




The five non-bipartite GRINs



The five bipartite GRINs which are star-like trees



The four bipartite GRINs which are trees but not star-like



The four bipartite GRINs which are not trees

4.2 When GRINs can coexist

4.2.1 Theorem

Given any set of GRINs there exists a Salem graph of which all are induced subgraphs if and only if all are bipartite or all are not.

Proof

Recall first that a Salem graph is defined in such a way that it is allowed to have an eigenvalue less than -2 only if it is bipartite. This is because the eigenvalues of a bipartite graph are symmetrical about zero [Bi p53] so an eigenvalue less than -2 is unavoidable unless we exclude all bipartite graphs from being Salem, which we don't wish to do.

A Salem graph cannot contain both a bipartite and a non-bipartite GRIN as induced subgraphs for this reason:

If the Salem graph contains an induced non-bipartite graph as a subgraph then the graph itself is not bipartite. This means it has no eigenvalue less than -2 . But growing it from the induced bipartite GRIN subgraph means by interlacing that it must have an eigenvalue less than -2 . This is impossible¹.

The fact that all other possibilities can exist is proved by exhibiting them.


For each pair of bipartite GRINs and for each pair of non-bipartite GRINs, here (in sections 4.2.4 and 4.2.5) are the smallest Salem graphs which contain both as induced subgraphs. "Smallest" means having the fewest vertices, and within that constraint having the fewest edges. Where there is more than one such graph, all are exhibited. Single thicker lines show the edges of one induced subgraph; double thinner lines show the edges of the other.

¹ The theorem was first observed after very extensive listing of graphs and their pairs of induced subgraphs, and it was proved (in this direction) by an exhaustive computer method. This was done at an early stage, before the problems with PARI's file handling had been resolved, and required three computers with a great deal of data transfer between them, working continuously for almost a month. This work was carried out with tenacious, determined enthusiasm, which in retrospect is probably better described as blind obstinate stupidity. The elation and exhilaration of triumph was as euphoric as it was misguided; the (not difficult) three-line proof given above emerged a couple of weeks later.

4.2.2 Different configurations

The induced subgraphs have been indicated in every case so that the reader can verify with minimal effort that they do exist as claimed.


This indication could often be made in more than one way. The graph below has B4 and B12 as induced subgraphs, and B12 can be highlighted in three different ways as shown. In the following pages, when such choice exists, the induced subgraphs shown have been chosen at random, because the question of when multiple copies of the same induced subgraph can exist within the same Salem graph is addressed extremely briefly in section 4.3.



4.2.3 A special edge







It is remarkable that in every case but one on the following pages, every edge in every graph belongs to at least one of the induced subgraphs. The sole exception is shown below.


This graph has A3 and A5 as induced subgraphs, with a dotted edge which does not belong to either subgraph. It connects (of course) part of the graph which belongs to only one of the induced subgraphs to part of the graph which belongs to only the other induced subgraph.





The smallest Salem graphs containing each possible pair of GRINs

4.2.4 The ten non-bipartite pairs






A3 A5 (2)




A4 A5


4.2.5 The 78 bipartite pairs




B1 B2




B1 B3




B1 B4




B1 B5 (2)








B1 B6 (2)




B1 B7




B1 B8




B1 B9 (5)




B1 B10




B1 B11



B1 B12




B1 B13




B2 B3




B2 B4




B2 B5







B2 B6 (2)





B2 B7 (2)






B2 B8 (2)



B2 B9 (2)




B2 B10 (3)













B2 B11








B2 B12




B2 B13




B3 B10 (5)




B3 B12 (3)



B4 B5 (4)




B4 B6 (2)




B4 B7 (2)




B4 B8




B4 B9 (3)










B4 B10 (8)




B4 B11




B4 B13 (2)







B5 B12 (6)




B5 B13 (4)




B6 B7 (2)




B6 B8 (3)




B6 B9 (3)




B6 B10 (6)




B6 B11 (2)




B6 B12




B6 B13




B7 B8 (2)




B7 B9




B7 B10 (7)




B7 B11 (2)




B7 B13




B7 B12 (3)




B8 B9 (2)




B8 B10 (8)




B8 B11 (2)




B8 B13




B8 B12 (4)




B9 B10 (9)




B9 B11 (2)




B9 B12 (5)




B9 B13 (3)




B10 B11 (2)




B10 B12 (5)




B10 B13 (6)



B11 B12



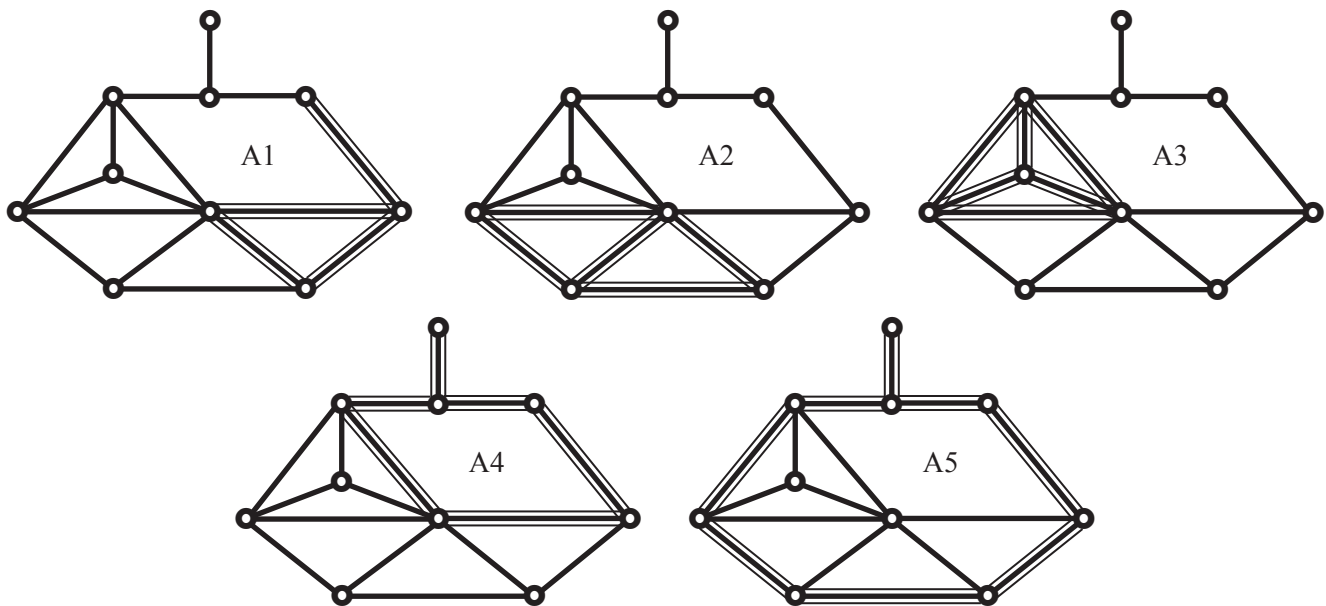
B11 B13



B12 B13 (4)

4.2.6 A Salem graph with all five non-bipartite GRINs as induced subgraphs


There are no Salem graphs with fewer than ten vertices which have all five non-bipartite GRINs as induced subgraphs. Remarkably, however, there are a 110 such graphs with ten vertices. These have between 16 and 22 edges. One of the 14 of these graphs with just 16 edges is shown below. It is illustrated five times, with each of the five non-bipartite GRINs shown in a separate example.



4.2.7 A Salem graph with all thirteen bipartite GRINs as induced subgraphs


There are no Salem graphs with fewer than thirteen vertices which have all the bipartite GRINs as induced subgraphs and there is, so far, no complete catalogue of Salem graphs with thirteen vertices. However, a Salem graph with thirteen vertices (which is only three more vertices than the largest GRIN) has been found. It has eighteen edges and is drawn below thirteen times, each showing a different bipartite GRIN as an induced subgraph.

It is not yet known whether there are other such graphs with thirteen vertices, possibly with fewer edges.



4.3 Minimal non-Salem graphs

As a very brief introduction to a different sort of minimal graph, here is a minimal non-Salem graph. Although not a Salem graph itself, the removal of any vertex, along with its associated edges, leaves a Salem graph.



Chapter 5

Other small Salem graphs of interest

Many of the graphs in this chapter will have been considered before, but not in a systematic study of small Salem graphs. All enumerations and complete sets of graphs are original work (except 5.2.3), as are the attempts to present the structures of some larger graphs in a memorable way.

5.1 Complementary graphs

5.1.1 Introduction


This section considers small connected Salem graphs with complements which are also connected Salem graphs. Those which most clearly fulfill this criterion are those which are their own complements. Since a graph and its complement must each have exactly half of all the possible edges, a self-complementary graph with n vertices must have $\frac{1}{4}n(n-1)$ edges. Either n or $(n-1)$ must be odd, so the other must be divisible by four. Hence there cannot exist self-complementary graphs with 3, 6, 7, 10 or 11 vertices.

Table 5.1 The numbers of complementary graphs with fewer than twelve vertices



Vertices	Self-complementary	Complementary in pairs	Total
3	Impossible	0	0
4	0	0	0
5	1	2	3
6	Impossible	30	30
7	Impossible	74	74
8	3	114	117
9	5	94	99
10	Impossible	74	74
11	Impossible	34	34

All these self-complementary graphs are now drawn, as well as all those with fewer than seven vertices. The second member of the pair is drawn twice: once to show clearly that it is the complement of the first, and secondly to show it rearranged into a more memorable shape.




5.1.2 The results for five vertices






5.1.3 The 15 complementary pairs with 6 vertices






$\bar{\text{I}}$






$\bar{\text{I}}$




$\bar{\text{I}}$





$\bar{\text{I}}$






$\bar{\text{I}}$











\mathbb{R}






\mathbb{R}




\mathbb{R}





\mathbb{R}






\mathbb{R}











\mathbb{R}






\mathbb{R}




\mathbb{R}



\mathbb{R}




\mathbb{R}




5.1.4 The 3 self-complementary graphs with 8 vertices

The first is also integral with spectrum $\{4, 1^{(2)}, 0, -1^{(2)}, -2^{(2)}\}$



5.1.5 The 5 self-complementary graphs with 9 vertices

The first graph is also integral with spectrum $\{4, 1^{(4)}, -2^{(4)}\}$



The last case is a wonderful example of how the salient properties of a graph can leap off the page if one is lucky enough to stumble upon a good way of drawing it.

When constructing the complementary graph the pentagon at the top turns into an isomorphic pentagram, and the vertices which are connected to all or none of those of the pentagon/pentagram are swapped (and are no longer joined), whereas the two vertices which are initially adjacent to a single edge become connected to each other, with each connected to the other vertex which is newly adjacent to a single edge.

5.2 Cospectral graphs

5.2.1 Introduction

This section considers sets of small connected Salem graphs which share the same characteristic equation and thus the same spectrum of eigenvalues. Most of these sets have two members and are commonly called PINGs (Pairs of Isospectral Non-isomorphic Graphs). In graphs with eleven vertices we also find sets with up to and including ten graphs.

Table 5.2 Numbers of sets of isospectral Salem graphs with fewer than twelve vertices

Vertices	6	7	8	9	10	11
Sets of 2 graphs (PINGs)	1	17	104	494	2002	6590
Sets of 3 graphs			11	68	384	1457
Sets of 4 graphs			1	10	85	463
Sets of 5 graphs				1	14	144
Sets of 6 graphs					5	57
Sets of 7 graphs					1	15
Sets of 8 graphs						7
Sets of 9 graphs						4
Sets of 10 graphs						2

5.2.2 Six vertices

There is a unique Salem PING with fewer than seven vertices:



Their characteristic equation is









































$$x^6 - 7x^4 - x^3 + 7x^2 + 4x - 1.$$

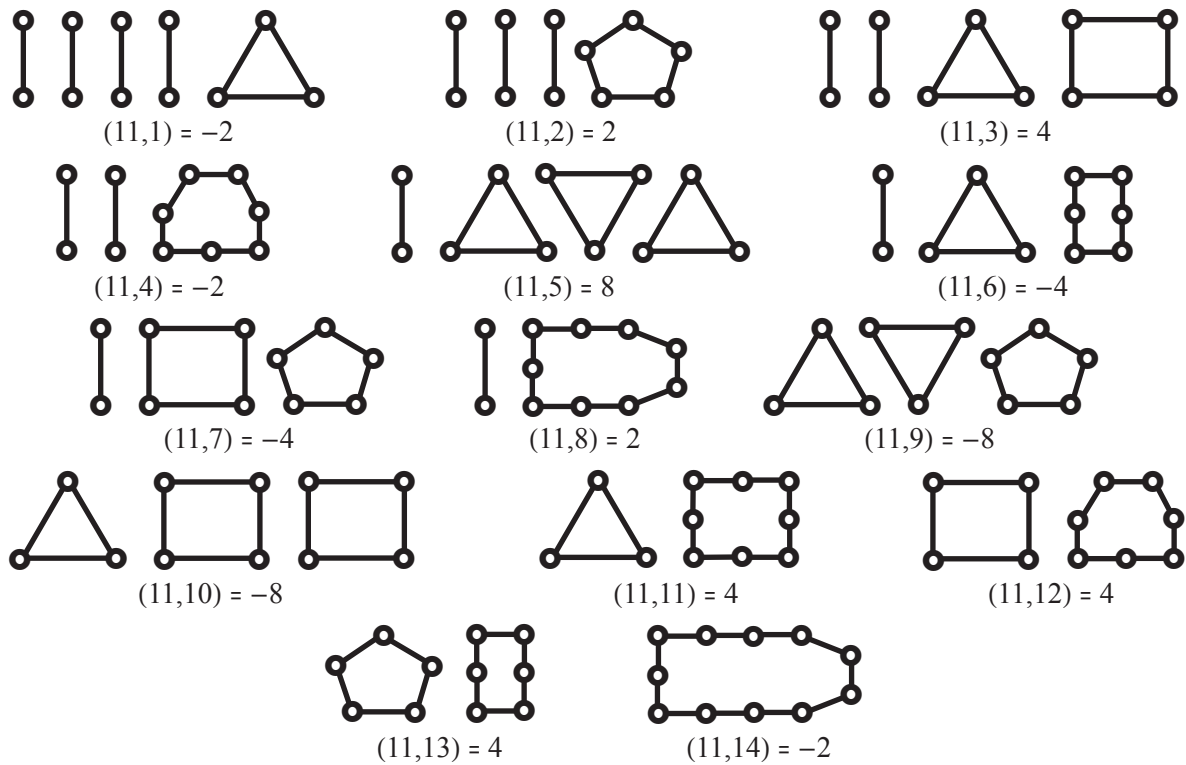
5.2.3 Elementary graphs

Illustrating these graphs gives an opportunity to describe a very interesting relationship between the prevalence of certain subgraphs of a graph, called elementary graphs, and its characteristic equation.

An elementary graph is frequently not connected. Each of its components is either a copy of K_2 or one of the cyclic graphs from C_3 upwards. With each graph the number $(-1)^c 2^k$ is associated, where c is the number of components, and k is the number of components which are **not** K_2 . Clearly k is a non-negative integer.

Here are all the elementary graphs with fewer than 12 vertices. The notation $(10,6) = -4$, for example, means that the sixth graph in the list of those with ten vertices has the associated number -4 .





The elementary graphs for each new vertex number are easily constructed from those already found, adding a new cycle (or K_2 initially) at each stage. Alternatively they can be seen as being in correspondence with those partitions of the integer representing the number of vertices which do not include the number one.

Let the characteristic equation of the graph under consideration be

$$c_0x^n + c_1x^{n-1} + c_2x^{n-2} + \dots + c_n.$$

Then c_0 is always one, c_1 is always zero, and to calculate the other coefficients c_i it is only necessary to count the occurrences of each elementary subgraph with i vertices, multiply by the associated number, and sum. This is proved in [B Proposition 7.3], using only the standard expansion of the determinant by minors and some very simple properties of adjacency matrices.



The prevalence of elementary subgraphs is “rather loosely” [Bi p46] related to the structure of the parent graph, and this, together with the fact that the associated numbers are always of the limited form $\pm 2^n$ where n is a non-negative integer, probably explains why there are fewer distinct characteristic polynomials (and so more PINGs) than one might at first suppose.

Let us return to the first example, now with the edges labelled (on the next page), and calculate the characteristic equation.

c_0 This is one as always.

c_1 This is zero as always.

c_2 $(2,1) = -1$ and in each graph there are seven edges, so $c_2 = 7(2,1) = -7$.



c_3 $(3,1) = -2$ and in each graph there are two triangles, so $c_3 = 2(3,1) = -4$.

c_4 This is the most interesting coefficient in this example because it demonstrates how the same value can arise in different ways.

$(4,1) = 1$, and in the first graph there are nine copies of $(4,1)$, specifically 1-4, 1-5, 1-6, 1-7, 2-6, 2-7, 3-5, 3-7 and 4-7. Also $(4,2) = -2$ and there is one copy of $(4,2)$, specifically 2-3-5-6.

$$c_3 = 9(4,1) + 1(4,2) = 7 \text{ for the first graph.}$$

In the second graph there are 7 copies of $(4,1)$, specifically a-d, a-e, a-f, a-g, b-g, c-g, and d-g.

$$c_3 = 7(4,1) = 7 \text{ for the second graph.}$$

c_5 $(5,1) = 2$ and the graph $(5,1)$ appears twice in each case, as 1-4-5-6, 2-3-4-7, a-b-c-g and a-e-f-g, so $c_5 = 2(5,1) = 4$ in both cases.

c_6 $(6,1) = -1$ and the graph $(6,1)$ appears once in each case, as 1-4-7 and a-d-g, so $c_6 = 1(6,1) = -1$ in both cases.

All this leads to the required characteristic equation:


$$x^6 - 7x^4 - x^3 + 7x^2 + 4x - 1.$$

At the very least this would seem to be an amusing party trick, enabling a characteristic equation to be conjured up using no algebra and minimal arithmetic. However, even with a highly systematic approach and the utmost concentration, the accurate counting of subgraphs is a matter of the very greatest difficulty. Except in the simplest cases one always ends up using the characteristic equation to check the counting rather than using the counting to construct the equation.

We now consider four sets of cospectral Salem graphs with eight vertices, and continue to calculate the characteristic equations since they are pivotal to the relationship between Salem numbers and graphs.

5.2.4 The smallest bipartite Salem PING

The smallest bipartite cospectral Salem graphs are eight pairs, each with eight vertices and with between seven and ten edges. The only pair with seven edges consists of these two trees:




The only elementary subgraphs which can be subsets of a tree are (2,1), (4,1), (6,1), (8,1) and so on, with associated numbers alternating between ± 1 , which goes a long way to explaining the relative simplicity of the characteristic equations of bipartate graphs and their relatively high number of PINGs.

In this case it is easy to see that there are no copies of (6,1) or (8,1) in either graph, and the copies of (2,1) are just the edges. It is not difficult to count nine copies of (4,1) in each case, so the characteristic equation is:

$$x^8 - 7x^6 + 9x^4.$$

5.2.5 The smallest set of three cospectral Salem graphs

The smallest sets of three cospectral Salem graphs have eight vertices. There are eleven sets, with between 10 and 18 edges. Only one set has 10 edges:




Characteristic equation:

$$x^8 - 10x^6 - 4x^5 + 25x^4 + 12x^3 - 16x^2 - 8x.$$

Coefficient	Graphs	Subgraphs	Result
c_0	A, B and C	None	1
c_1	A, B and C	None	0
c_2	A, B and C	$10(2,1) = -10$	-10
c_3	A, B and C	$2(3,1) = -4$	-4
c_4	A, B and C	$27(4,1) = 27$ $1(4,2) = -2$	25
c_5	A B C	$8(5,1) = 16$ $2(5,2) = -4$ $7(5,1) = 14$ $1(5,2) = -2$ $6(5,1) = 12$	12
c_6	A B and C	$18(6,1) = -18$ $3(6,2) = 6$ $2(6,4) = -4$ $20(6,1) = -20$ $3(6,2) = 6$ $1(6,4) = -2$	-16
c_7	A B C	$4(7,1) = -8$ $4(7,1) = -8$ $1(7,2) = 2$ $1(7,4) = -2$ $2(7,1) = -4$ $2(7,4) = -4$	-8
c_8	A, B and C	$2(8,1) = 2$ $1(8,2) = -2$	0

5.2.6 The smallest set of four cospectral Salem graphs

The members of this set have eight vertices and twelve edges:



Characteristic equation:


$$x^8 - 12x^6 - 10x^5 + 28x^4 + 28x^3 - 15x^2 - 12x + 4.$$

Coefficient	Graphs	Subgraphs	Result
c_0	A, B, C and D	None	1
c_1	A, B, C and D	None	0
c_2	A, B, C and D	$12(2,1) = -12$	-12
c_3	A, B, C and D	$5(3,1) = -10$	-10
c_4	A	$34(4,1) = 34$ $3(4,2) = -6$	28
	B	$38(4,1) = 38$ $5(4,2) = -10$	
	C and D	$36(4,1) = 36$ $4(4,2) = -8$	
c_5	A, B and D	$17(5,1) = 34$ $3(5,2) = -6$	28
	C	$18(5,1) = 36$ $4(5,2) = -8$	
c_6	A	$23(6,1) = -23$ $3(6,2) = 6$ $1(6,2) = 4$ $1(6,4) = -2$	-15
	B	$35(6,1) = -35$ $12(6,2) = 24$ $2(6,4) = -4$	
	C	$27(6,1) = -27$ $7(6,2) = 14$ $1(6,2) = 4$ $3(6,4) = -6$	
	D	$27(6,1) = -27$ $8(6,2) = 16$ $1(6,2) = 4$ $4(6,4) = -8$	
c_7	A	$9(7,1) = -18$ $3(7,2) = 6$	-12
	B	$11(7,1) = -22$ $7(7,2) = 14$ $2(7,4) = -4$	
	C	$9(7,1) = -18$ $4(7,2) = 8$ $1(7,4) = -2$	
	D	$7(7,1) = -14$ $3(7,2) = 6$ $2(7,4) = -4$	
c_8	A, C, and D	$2(8,1) = 2$ $1(8,4) = 2$	4
	B	$4(8,1) = 4$ $2(8,2) = -4$ $2(8,4) = 4$ $1(8,5) = 4$ $2(8,7) = -4$	

5.2.7 The smallest integral cospectral Salem graphs

These are in a set of three, with eight vertices and fourteen edges. The characteristic equation and spectrum are:

$$x^8 - 14x^6 - 16x^5 + 25x^4 + 32x^3 - 12x^2 - 16x \text{ and } \{4, 1^{(2)}, 0, -1^{(2)}, -2^{(2)}\}.$$

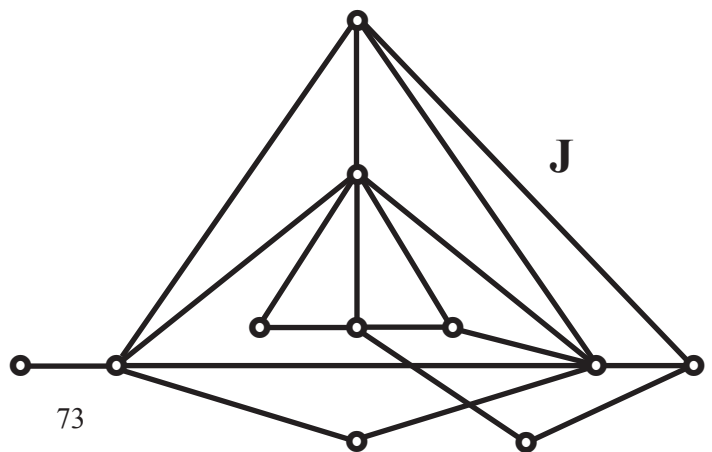











Coefficient	Graphs	Subgraphs	Result
c_0	A, B and C	None	1
c_1	A, B and C	None	0
c_2	A, B and C	$14(2,1) = -14$	-14
c_3	A, Btt and C	$8(3,1) = -16$	-16
c_4	A and C B	$47(4,1) = 47$ $11(4,2) = -22$ $49(4,1) = 49$ $12(4,2) = -24$	25
c_5	A and C B	$28(5,1) = 56$ $12(5,2) = -24$ $31(5,1) = 62$ $15(5,2) = -30$	32
c_6	A B C	$36(6,1) = -36$ $16(6,2) = 32$ $2(6,3) = 8$ $8(6,4) = -16$ $42(6,1) = -42$ $27(6,2) = 54$ $1(6,3) = 4$ $14(6,4) = -28$ $40(6,1) = -40$ $18(6,2) = 36$ $2(6,3) = 8$ $8(6,4) = -16$	-12
c_7	A B C	$12(7,1) = -24$ $8(7,2) = 16$ $4(7,4) = -8$ $15(7,1) = -30$ $15(7,2) = 30$ $1(7,3) = 4$ $10(7,4) = -20$ $20(7,1) = -40$ $12(7,2) = 24$	-16
c_8	A B C	$2(8,1) = 2$ $1(8,7) = -2$ $4(8,1) = 4$ $4(8,2) = -8$ $2(8,4) = 4$ $6(8,1) = 6$ $7(8,2) = -14$ $2(8,3) = -8$ $8(8,4) = 16$	0

5.2.8 The smallest set of ten cospectral Salem graphs

There are two sets of ten cospectral Salem graphs with eleven vertices, one with 20 edges, and the other, considered here, with 19 edges. The characteristic equation is:

$$x^{11} - 19x^9 - 18x^8 + 88x^7 + 108x^6 - 127x^5 - 156x^4 + 64x^3 + 48x^2 - 16x.$$



Coefficient	Graphs	Subgraphs				Result
c_0	All	None				1
c_1	All	None				0
c_2	All	$19(2,1) = -19$				-19
c_3	All	$9(3,1) = -18$				-18
c_4	A, F, I B, C, D, E, G, H, J	$112(4,1) = 112$ $110(4,1) = 110$	$12(4,2) = -24$ $11(4,2) = -22$			88
c_5	A	$74(5,1) = 148$	$20(5,2) = -40$			108
	B	$68(5,1) = 136$	$14(5,2) = -28$			
	C, J	$69(5,1) = 138$	$15(5,2) = -30$			
	D, E, H	$70(5,1) = 140$	$16(5,2) = -32$			
	F, I	$72(5,1) = 144$	$18(5,2) = -36$			
	J	$71(5,1) = 142$	$17(5,2) = -34$			
c_6	A	$243(6,1) = -243$	$69(6,2) = 138$	$8(6,3) = 32$	$27(6,4) = -54$	-127
	B	$237(6,1) = -237$	$60(6,2) = 120$	$5(6,3) = 20$	$15(6,4) = -30$	
	C	$235(6,1) = -236$	$61(6,2) = 122$	$5(6,3) = 20$	$17(6,4) = -34$	
	D	$231(6,1) = -231$	$60(6,2) = 120$	$7(6,3) = 28$	$22(6,4) = -44$	
	E	$231(6,1) = -231$	$58(6,2) = 116$	$7(6,3) = 28$	$20(6,4) = -40$	
	F	$247(6,1) = -247$	$64(6,2) = 128$	$7(6,3) = 28$	$18(6,4) = -36$	
	G	$227(6,1) = -227$	$61(6,2) = 122$	$7(6,3) = 28$	$25(6,4) = -50$	
	H	$229(6,1) = -229$	$60(6,2) = 120$	$6(6,3) = 24$	$21(6,4) = -42$	
	I	$243(6,1) = -243$	$65(6,2) = 130$	$9(6,3) = 36$	$25(6,4) = -50$	
	J	$229(6,1) = -229$	$58(6,2) = 116$	$6(6,3) = 24$	$19(6,4) = -38$	
c_7	A	$153(7,1) = -306$	$85(7,2) = 170$	$9(7,3) = 36$	$28(7,4) = -56$	-156
	B	$136(7,1) = -272$	$54(7,2) = 108$	$7(7,3) = 28$	$10(7,4) = -20$	
	C	$139(7,1) = -278$	$60(7,2) = 120$	$8(7,3) = 32$	$15(7,4) = -30$	
	D	$130(7,1) = -260$	$56(7,2) = 112$	$8(7,3) = 32$	$20(7,4) = -40$	
	E	$138(7,1) = -276$	$63(7,2) = 126$	$8(7,3) = 32$	$19(7,4) = -38$	
	F	$155(7,1) = -312$	$76(7,2) = 152$	$7(7,3) = 28$	$12(7,4) = -24$	
	G	$130(7,1) = -260$	$64(7,2) = 128$	$7(7,3) = 28$	$26(7,4) = -52$	
	H	$131(7,1) = -262$	$61(7,2) = 122$	$6(7,3) = 24$	$20(7,4) = -40$	
	I	$147(7,1) = -294$	$79(7,2) = 158$	$12(7,3) = 48$	$34(7,4) = -68$	
	J	$127(7,1) = -254$	$58(7,2) = 116$	$7(7,3) = 28$	$23(7,4) = -46$	
c_8	A {	$174(8,1) = 174$	$84(8,2) = -168$	$21(8,3) = -84$	$73(8,4) = 146$	64
		$11(8,5) = 44$	$1(8,6) = 4$	$26(8,7) = -52$		
	B {	$180(8,1) = 180$	$82(8,2) = -164$	$12(8,3) = -48$	$36(8,4) = 72$	
		$5(8,5) = 20$	$5(8,6) = 20$	$8(8,7) = -16$		
	C {	$172(8,1) = 172$	$82(8,2) = -164$	$12(8,3) = -48$	$42(8,4) = 84$	
		$5(8,5) = 20$	$6(8,6) = 24$	$12(8,7) = -24$		
	D {	$168(8,1) = 168$	$75(8,2) = -150$	$12(8,3) = -48$	$44(8,4) = 88$	
		$5(8,5) = 20$	$2(8,6) = 8$	$11(8,7) = -22$		
	E {	$160(8,1) = 160$	$70(8,2) = -140$	$16(8,3) = -64$	$49(8,4) = 98$	
		$9(8,5) = 36$	$1(8,6) = 4$	$15(8,7) = -30$		
	F {	$182(8,1) = 182$	$80(8,2) = -160$	$23(8,3) = -92$		
		$61(8,4) = 122$	$11(8,5) = 44$	$16(8,7) = -32$		
G {	$146(8,1) = 146$	$67(8,2) = -134$	$11(8,3) = -44$	$58(8,4) = 116$		
	$3(8,5) = 12$	$2(8,6) = 8$	$20(8,7) = -40$			
H {	$152(8,1) = 152$	$70(8,2) = -140$	$11(8,3) = -44$	$47(8,4) = 94$		
	$5(8,5) = 20$	$2(8,6) = 8$	$13(8,7) = -26$			
I {	$170(8,1) = 170$	$67(8,2) = -134$	$25(8,3) = -100$			
	$70(8,4) = 140$	$14(8,5) = 56$	$34(8,7) = -68$			
J {	$150(8,1) = 150$	$62(8,2) = -124$	$10(8,3) = -40$	$47(8,4) = 94$		
	$4(8,5) = 16$	$2(8,6) = 8$	$20(8,7) = -40$			

$$\begin{array}{l}
c_9 \\
A \left\{ \begin{array}{l} 174(9,1)=174 \quad 84(9,2)=-168 \quad 21(9,3)=-84 \quad 73(9,4)=146 \\ 11(9,5)=-44 \quad 1(9,6)=4 \quad 26(9,7)=52 \quad 26(9,8)=-52 \end{array} \right. \\
B \left\{ \begin{array}{l} 180(9,1)=180 \quad 82(9,2)=-164 \quad 12(9,3)=-48 \\ 5(9,4)=20 \quad 5(9,7)=20 \quad 8(9,8)=-16 \end{array} \right. \\
C \left\{ \begin{array}{l} 172(9,1)=172 \quad 82(9,2)=-164 \quad 12(9,3)=-48 \\ 5(9,4)=20 \quad 6(9,7)=24 \quad 12(9,8)=-24 \end{array} \right. \\
D \left\{ \begin{array}{l} 168(9,1)=168 \quad 75(9,2)=-150 \quad 12(9,3)=-48 \quad 44(9,4)=88 \\ 5(9,6)=-20 \quad 2(9,7)=8 \quad 11(9,8)=-22 \end{array} \right. \\
E \left\{ \begin{array}{l} 160(9,1)=160 \quad 70(9,2)=-140 \quad 16(9,3)=-64 \quad 49(9,4)=98 \\ 9(9,6)=36 \quad 1(9,7)=4 \quad 15(9,8)=-30 \end{array} \right. \\
F \left\{ \begin{array}{l} 182(9,1)=182 \quad 80(9,2)=-160 \quad 23(9,3)=-92 \quad 23(9,4)=92 \\ 61(9,5)=122 \quad 11(9,6)=44 \quad 16(9,7)=32 \quad 23(9,8)=-92 \end{array} \right. \\
G \left\{ \begin{array}{l} 146(9,1)=146 \quad 67(9,2)=-134 \quad 11(9,3)=-44 \quad 58(9,4)=116 \\ 3(9,6)=12 \quad 2(9,7)=8 \quad 20(9,8)=-40 \end{array} \right. \\
H \left\{ \begin{array}{l} 152(9,1)=152 \quad 70(9,2)=-140 \quad 11(9,3)=-44 \quad 47(9,4)=94 \\ 5(9,6)=-20 \quad 2(9,7)=8 \quad 13(9,8)=-26 \end{array} \right. \\
I \left\{ \begin{array}{l} 170(9,1)=170 \quad 67(9,2)=-134 \quad 25(9,3)=-100 \quad 23(9,4)=92 \\ 70(9,5)=-140 \quad 14(9,6)=56 \quad 34(9,8)=-68 \end{array} \right. \\
J \left\{ \begin{array}{l} 150(9,1)=150 \quad 62(9,2)=-124 \quad 10(9,3)=-40 \quad 47(9,4)=94 \\ 4(9,6)=16 \quad 2(9,7)=8 \quad 20(9,8)=-40 \end{array} \right.
\end{array}$$

48

$$\begin{array}{l}
c_{10} \\
A \left\{ \begin{array}{l} 28(10,1)=-28 \quad 18(10,2)=36 \quad (10,3)=12 \quad 24(10,4)=-48 \\ 3(10,5)=-12 \quad 1(10,6)=-4 \quad 14(10,7)=28 \end{array} \right. \\
B \left\{ \begin{array}{l} 34(10,1)=-34 \quad 25(10,2)=50 \quad 6(10,3)=24 \quad 14(10,4)=-28 \\ 5(10,5)=-20 \quad 4(10,6)=-16 \quad 6(10,7)=12 \quad 1(10,8)=-8 \\ 1(10,10)=4 \quad 1(10,11)=4 \quad 2(10,12)=-4 \end{array} \right. \\
C \left\{ \begin{array}{l} 30(10,1)=-30 \quad 23(10,2)=46 \quad 6(10,3)=24 \quad 16(10,4)=-32 \\ 5(10,5)=-20 \quad 4(10,6)=-16 \quad 6(10,7)=12 \quad 1(10,8)=-8 \\ 2(10,10)=8 \quad 1(10,11)=4 \quad 2(10,12)=-4 \end{array} \right. \\
D \left\{ \begin{array}{l} 30(10,1)=-30 \quad 20(10,2)=40 \quad 3(10,3)=12 \quad 14(10,4)=-28 \\ 4(10,5)=-16 \quad 2(10,6)=-8 \quad 5(10,7)=10 \quad 1(10,11)=4 \end{array} \right. \\
E \left\{ \begin{array}{l} 26(10,1)=-26 \quad 16(10,2)=32 \quad 5(10,3)=20 \quad 14(10,4)=-28 \\ 5(10,5)=-20 \quad 1(10,6)=-4 \quad 7(10,7)=14 \\ 1(10,8)=-8 \quad 2(10,9)=8 \quad 2(10,12)=-4 \end{array} \right. \\
F \left\{ \begin{array}{l} 30(10,1)=-30 \quad 17(10,2)=34 \quad 7(10,3)=28 \\ 20(10,4)=-40 \quad 5(10,5)=-20 \quad 8(10,7)=16 \\ 2(10,9)=8 \quad 2(10,10)=8 \quad 10(10,12)=-20 \end{array} \right. \\
G \quad 16(10,1)=-16 \quad 6(10,2)=12 \quad 8(10,4)=-16 \quad 2(10,7)=4 \\
H \left\{ \begin{array}{l} 18(10,1)=-18 \quad 9(10,2)=18 \quad 9(10,4)=-18 \\ 1(10,6)=-4 \quad 4(10,7)=8 \quad 1(10,12)=-2 \end{array} \right. \\
I \left\{ \begin{array}{l} 26(10,1)=-26 \quad 9(10,2)=18 \quad 5(10,3)=20 \\ 20(10,4)=-40 \quad 3(10,5)=-12 \quad 12(10,7)=24 \end{array} \right. \\
J \left\{ \begin{array}{l} 18(10,1)=-18 \quad 9(10,2)=18 \quad 9(10,4)=-18 \\ 1(10,6)=-4 \quad 4(10,7)=8 \quad 1(10,12)=-2 \end{array} \right.
\end{array}$$

-16

$$\begin{array}{l}
c_{11} \\
A, I \quad 2(11,1)=-4 \quad 4(11,2)=8 \quad 2(11,4)=-4 \\
B \left\{ \begin{array}{l} 8(11,1)=-16 \quad 6(11,2)=12 \quad 3(11,3)=12 \\ 2(11,4)=-4 \quad 2(11,7)=-8 \quad 2(11,8)=4 \end{array} \right. \\
C \quad 6(11,1)=-12 \quad 4(11,2)=8 \quad 2(11,3)=8 \quad 2(11,4)=-4 \\
D \left\{ \begin{array}{l} 6(11,1)=-12 \quad 6(11,2)=12 \quad 2(11,3)=8 \\ 2(11,4)=-4 \quad 1(11,7)=-4 \end{array} \right. \\
E \quad 4(11,1)=-8 \quad 4(11,2)=8 \quad 1(11,3)=4 \quad 2(11,4)=-4 \\
F \left\{ \begin{array}{l} 4(11,1)=-8 \quad 2(11,2)=4 \quad 1(11,3)=4 \\ 2(11,4)=-4 \quad 2(11,8)=4 \end{array} \right. \\
G, H, J \quad \text{None}
\end{array}$$

0

5.3 Integral and trivial Salem graphs


5.3.1 Introduction

This section considers small connected Salem graphs with all their eigenvalues integral, and those with an integer as their greatest eigenvalue. A table of the numbers of each with fewer than twelve vertices, broken down by number of edges, is given on the next page. All the integral examples with fewer than nine vertices are drawn, a few infinite sets and well-known larger examples are described, and finally the smallest trivial examples are considered.


5.3.2 The smallest examples

The smallest integral Salem graph is K_4


There are two examples with five vertices, including K_5



$$\{3, -1^{(3)}\}$$




$$\{3, 0^{(2)}, -1, -2\}$$




$$\{4, -1^{(4)}\}$$


The four examples with six vertices are: the triangular prism, which can also be regarded as a rather trivial example of an extended Petersen graph $GP(3,1)$; the utility graph $K_{3,3}$; the cocktail party graph which is complete apart from $\frac{1}{2}n$ independent edges (where $n=6$ here); and the complete graph K_6 .




$$\{3, 1, 0^{(2)}, -2^{(2)}\}$$



$$\{\pm 3, 0^{(4)}\}$$



$$\{4, 0^{(3)}, -2^{(2)}\}$$




$$\{5, -1^{(5)}\}$$


There are five examples with seven vertices.

The friendship graphs (also known as Dutch windmill graphs), based on this pattern, are the only finite graphs in which every pair of vertices has exactly one neighbour [ERS]. The smallest, C_3 (a single sail), is cyclotomic; the next, (also known as the bow tie graph or butterfly graph), is Salem but not integral.


No other friendship graph is Salem.




$$\{3, 1^{(2)}, -1^{(3)}, -2\}$$




$$\{3, 2, 0, -1^{(3)}, -2\}$$



$$\{3, 1^{(2)}, 0, -1, -2^{(2)}\}$$



$$\{5, 1, -1^{(4)}, -2\}$$





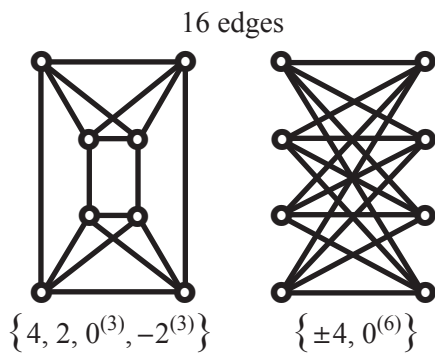










$$\{6, -1^{(6)}\}$$

Table 5.3 Numbers of integral Salem graphs and (in brackets) trivial Salem graphs with fewer than twelve vertices, listed according to their number of edges


Edges	Vertices							
	4	5	6	7	8	9	10	11
6	1							
7		1						
8		0						
9		0	2	1	1		1	
10		1	0	2	2 (2)		0	
11			0	0	2 (0)	(1)	0	
12			1	0	1 (0)	1 (4)	3 (2)	(1)
13			0	0	1 (0)	0 (2)	2 (5)	1 (4)
14			0	0	3 (0)	0 (0)	3 (3)	4 (8)
15			1	0	1 (1)	1 (1)	4 (2)	4 (9)
16				0	2 (1)	4 (5)	1 (2)	0 (0)
17				1	0	1 (2)	2 (3)	1 (0)
18				0	1	2 (1)	6 (4)	3 (1)
19				0	2	0 (0)	2 (2)	4 (5)
20				0	0	0 (0)	2 (0)	3 (2)
21				1	0	1 (0)	0 (0)	0 (3)
22					0	0 (0)	0 (0)	2 (0)
23					0	0 (0)	6 (1)	6 (3)
24					1	0 (0)	0 (0)	0 (3)
25					0	0 (1)	1 (0)	4 (3)
26					0	2 (0)	1 (0)	1 (1)
27					0	0 (0)	1 (1)	0 (0)
28					1	0 (0)	2 (2)	1 (1)
29						0 (0)	1 (1)	0 (2)
30						0 (0)	1	1 (5)
31						0 (1)	0	4 (1)
32						0	0	1 (1)
33						0	2	1 (0)
34						0	0	0 (0)
35						0	0	2 (0)
36						1	0	0 (1)
37							0	3 (1)
38							0	0 (0)
39							0	0 (0)
40							1	0 (0)
41							0	0 (0)
42							0	0 (0)
43							0	2 (1)
44							0	0
45							1	0
46								0
47								0
48								0
49								2
50								0
51								0
52								0
53								0
54								0
55								1
Integral:	1	2	4	5	18	13	43	51
Trivial:					(4)	(18)	(28)	(56)

5.3.3 The eighteen integral Salem graphs with eight vertices






5.3.4 A selection of infinite sets of integral Salem graphs and special named integral Salem graphs


The star graph with 10 vertices




The standard Petersen graph GP(5,2)



The Nauru¹ graph GP(12,5)



The line graph of $K_{4,4}$



5.3.5 The star graphs

It is easily seen in the table that with fewer than twelve vertices there is only one integral Salem tree (with number of edges one fewer than number of vertices). This is the star graph shown above. All star graphs are Salem, except those with fewer than six vertices which are cyclotomic. Those with n vertices, where $n-1$ is a perfect square, are also integral with spectrum $\{\pm\sqrt{n-1}, 0^{(n-2)}\}$.

5.3.6 The complete graphs

The complete graphs K_n are all integral with spectrum $\{n-1, -1^{(n-1)}\}$. K_1 to K_3 are cyclotomic, and the rest are Salem. K_4 to K_{11} are seen in the table at the bottom of each integral column.

¹ The Nauru graph was named by David Eppstein after the island nation, which has a twelve-pointed star on its flag representing the original twelve tribes of Nauru. Remarkably, if the [currently universally recognized] states of the world are arranged in alphabetical order, each of the three consecutive states Namibia, Nauru and Nepal (and no other) has a twelve-pointed star on its flag.

5.3.7 The complete bipartite graphs

The complete bipartite graph $K_{m,n}$ has spectrum $\{\pm\sqrt{mn}, 0^{(m+n-2)}\}$, so $K_{1,1}$, $K_{2,1}$, $K_{3,1}$, $K_{4,1}$, and $K_{2,2}$ are cyclotomic while the rest are Salem. The integral cases can be thought of as:

- $K_{i^2,1}$ ($i > 2$), the integral star graphs described above.
- $K_{i,i}$ ($i > 2$). $K_{3,3}$, $K_{4,4}$, and $K_{5,5}$ appear in the table; the first two are also drawn among the graphs with six and eight vertices.
- Other cases $K_{m,n}$ where mn is a perfect square, such as $K_{12,3}$.

5.3.8 The cocktail party graphs

The cocktail party graph CP_n (n even, ≤ 4) is the same as K_n with (any) $\frac{1}{2}n$ independent edges removed. This is frequently represented as the complete graph drawn as a regular polygon without the edges joining opposite vertices.

The spectrum is $\{n-2, 0^{(\frac{1}{2}n)}, -2^{(\frac{1}{2}n-1)}\}$, so CP_4 is cyclotomic (and isomorphic to C_4) while the rest are integral Salem. CP_6 , CP_8 and CP_{10} appear in the table as the penultimate entry in their respective columns; the first two are also drawn among the graphs with six and eight vertices.

5.3.9 The generalized Petersen graphs

The generalized Petersen graph $GP(n,k)$ is defined as follows:

The vertices are $u_0 \dots u_{n-1}$ and $v_0 \dots v_{n-1}$. These are often drawn as the n vertices of each of two concentric regular polygons.

The edges fall into three sets:

- $(u_i, u_{i+1 \pmod n})$. These complete the outer polygon in the obvious way.
- (u_i, v_i) . These are the “spokes” joining the inner and outer polygons.
- $(v_i, v_{i+k \pmod n})$. These complete the inner polygon, “winding it up”. In the standard Petersen graph for example, shown at the top of this section, the inner pentagon “winds up” into a pentagram. If $k=1$ there is no “winding” and we have the graph of a polygonal prism such as $GP(6,1)$ shown at the top of this section. If n and k are not co-prime the inner polygon degenerates into separate components; an extreme example being $GP(6,3)$, shown at the top of this section, in which the inner hexagon degenerates into three independent edges.

Eight of the smaller generalized Petersen graphs are Salem:

- $GP(3,1)$, the triangular prism, is integral Salem, illustrated among the graphs with six vertices. The spectrum is $\{3, 1, 0^{(2)}, -2^{(2)}\}$.
- $GP(4,1)$, the cube, is integral Salem, illustrated among the graphs with eight vertices. The spectrum is $\{\pm 3, \pm 1^{(3)}\}$.

- GP(6,1), the hexagonal prism, is integral Salem, illustrated at the top of this section. The spectrum is $\{\pm 3, \pm 2^{(2)}, \pm 1, 0^{(4)}\}$.
- GP(5,2), the standard Petersen graph, is integral Salem, illustrated at the top of this section. The spectrum is $\{3, 1^{(5)}, -2^{(4)}\}$.
- GP(10,3), the Desargues graph, is integral Salem with spectrum $\{\pm 3, \pm 2^{(4)}, \pm 1^{(5)}\}$.
- GP(12,5), the Nauru graph, is integral Salem, illustrated at the top of this section. The spectrum is $\{\pm 3, \pm 2^{(6)}, \pm 1^{(3)}, 0^{(4)}\}$.
- GP(8,3), the Möbius-Kantor graph, is trivial Salem, with the following spectrum: $\{\pm 3, \pm \sqrt{3}^{(4)}, \pm 1^{(3)}\}$.
- GP(6,3), illustrated at the top of this section, is a degenerate generalized Petersen graph, but a “proper” Salem graph.

5.3.10 The Shrikhande graph

The Shrikhande graph [Sh] has been singled out for special mention for a combination of reasons:

Most importantly it is an integral Salem graph. Its spectrum is $\{6, 2^{(6)}, -2^{(9)}\}$.

Secondly, it has already been encountered in §2.6.1, where it was most conveniently described as a triangulation of the sphere, and most easily drawn in two dimensions as a triangular tessellation with some vertices appearing at least twice at the edges. It is of course possible to draw it in two dimensions with each vertex appearing only once, but all such drawings belie its inherent simplicity.


Thirdly, it has a remarkable unique property. The spectra of the line graphs of the complete bipartite graphs $K_{i,i}$ are unique, with the single exception that the spectrum of the line graph of $K_{4,4}$ is the same as that of the Shrikhande graph. The line graph of $K_{4,4}$ is shown at the top of this section. The thick gray lines each represent four edges between the copies of $K_{4,4}$. These edges go between corresponding points (i.e. top-left to top-left and so on).



Finally the Shrikhande graph is also a strongly regular integral Salem graph.

5.3.11 The smallest Salem graphs for which the largest eigenvalue is an integer

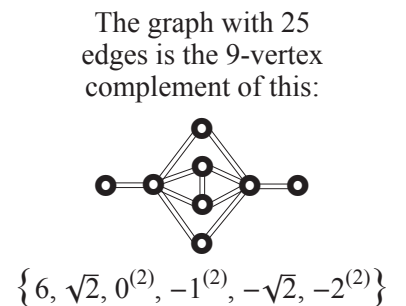
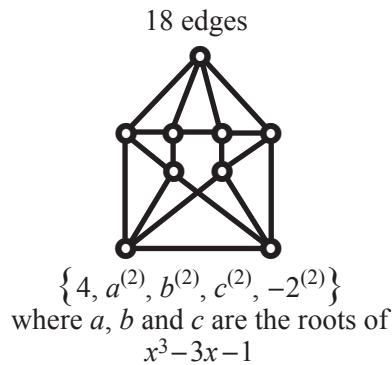
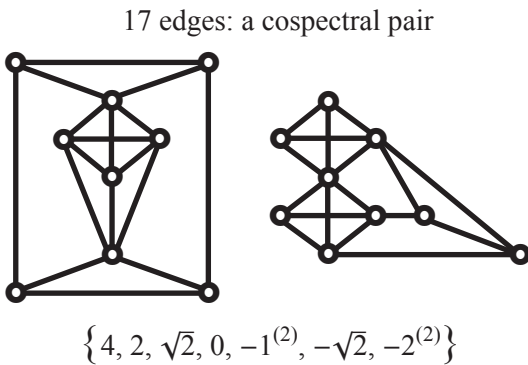










In all but one of these smallest cases all the non-integral eigenvalues are closely related to either $\sqrt{2}$, $\sqrt{3}$ or the golden ratio ϕ .

The four trivial Salem graphs with eight vertices:





The 18 trivial Salem graphs with nine vertices:



The last graph, with 31 edges, is the complete graph C_9 with a copy of C_5 removed.

Its spectrum is $\{7, \phi^{(2)}, -1^{(3)}, -\phi-1^{(2)}, -2\}$

5.4 Regular Salem graphs

5.4.1 Introduction

This section considers small connected regular Salem graphs, in particular those that satisfy the stronger requirements of being circulant or strongly regular, and including those that satisfy the weaker requirement of being semi-regular.

5.4.2 Circulant Salem graphs

A circulant graph has a circulant adjacency matrix which is defined by its first row; the entries of each subsequent row are the same as the entries of the previous row, but shifted one space to the right, with the one that “falls off the end” reappearing at the beginning.

If the first row of a circulant adjacency matrix is $(a_0 \ a_1 \ a_2 \ \dots \ a_{n-1})$ then $a_0=0$ because loops are excluded, and $a_i=a_{n-i}$ ($0 < i < n$) because the matrix must be symmetric as well as circulant.

This means that the number of circulant adjacency matrices is (relatively) very small. It is much easier to find those circulant graphs which are Salem and identify them with those regular Salem graphs already found, than it is to establish which of the regular Salem graphs already found are circulant. Also it is not difficult to push the exhaustive search for circulant Salem graphs up to 40 vertices; the results are shown in table 5.4.

The table very strongly suggests a complete classification of all circulant Salem graphs. There are five infinite sets, the most interesting of which depends on the prime factors of the number of vertices, and in addition there are seven small sporadic examples.


5.4.3 Remarks on drawing circulant graphs

Given several circulant adjacency matrices which are known to represent isomorphic graphs, then, when drawing the graph, we choose if possible one of the matrices in which $a_1=a_{n-1}=1$. This means that we can start the drawing with n vertices arranged in a regular polygon (numbered if necessary consecutively), with at least some of the edges between them completing the polygon in the obvious way, meaning that physically adjacent vertices are also adjacent in the graph. Such a choice is always possible in all of the infinite sets of graphs which we are about to consider; and a few cases where such a choice is not possible will be detailed among the sporadic examples.

If in addition n is even and $a_{\frac{1}{2}n-1}=1$ then the diagonals of this polygon can also be drawn. The polygon and its diagonals can then be redrawn in a way which is superficially very different, known as a Möbius ladder.

Consider the circulant adjacency matrices with first rows (0001010100) and (0100010001) . These represent isomorphic graphs. A way of seeing this with minimal effort will be described in section 5.4.6. We prefer the second matrix, and note that the diagonal edges are present. The graph can now be drawn in at least three ways.

The first is the circulant drawing which has the advantage of showing the symmetry beautifully.



The second is a Möbius ladder. This is vastly easier to draw, but the corner vertices are repeated and so need to be labelled to show that a twist has been introduced, making the (obvious) ladder similar to a Möbius strip.

The third drawing bends the ladder round into a circle. Vertices don't need to be repeated, the similarity to the graph of the (in this case) pentagonal prism is clear, and the twist is also self-evident; but the symmetry needs to be thought about — it no longer leaps off the page as in the circulant version — and the drawing once again takes some effort.

Now consider the graph with (01110111) as the first row of its adjacency matrix, which is in fact CP_8 , the cocktail party graph with eight vertices. Vertex 1 is connected to vertices 2, 3, 4, 6, 7 and 8, but it is much more informative to regard these as “jumps” of ± 1 (to vertices 2 and 8 in this case), of ± 2 (to vertices 3 and 7 in this case), and ± 3 (to vertices 4 and 6). These *relative* jumps, as opposed to the *absolute* numbers of the target vertices, apply to every vertex. The “ \pm ” pairs always occur because of the restrictions described in 4.4.1, except for the diagonals where it is not necessary because the result of the jump is the same regardless of the direction in which it is made. In this example jumps of +4 and -4 are the same, and represent the diagonals.

If the jump number does not divide the number of vertices, the jumps generate another polygon with all the vertices. This can either be drawn “wound up” inside the first polygon, or drawn separately. In the latter case the vertices must be labelled since their relationship to the vertices in the other polygon is a crucial part of the structure.

If the jump number does divide the number of vertices the second polygon arrives back at the start without including all the vertices so the new polygon decomposes into smaller polygons. The jump number is the number of smaller polygons and the quotient of the total number of vertices by the jump number is the number of vertices in the smaller polygons.


In this example the ± 1 and the ± 3 jumps generate full octagons, while the ± 2 jumps generate two squares. Some or all of these can then be combined together if clarity allows, as shown on the next page but one after the table.

In the previous example the jumps were ± 1 and 5, so the graph can (as is clear from the circulant drawing), be decomposed into a decagon and five copies of K_2 which are the diagonals.

Table 5.4 Circulant Salem graphs with fewer than 41 vertices, arranged by their n -regularity

n	Vertices																		
	4	6	8	10	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40
3	1	<u>2</u> *		1															
4		1	<u>2</u> *	1	2			<u>1</u>											
5		1	•	1	<u>1</u>	1		<u>1</u>											
6			1	•	1	1	1	•											
7			1	•	<u>1</u>	1	•	2											
8				1	•	•	1	1	2										
9				1	•	•	•	1	•	1									
10					1	•	•	•	1	1	2								
11					1	•	•	•	•	1	•	1							
12						1	•	•	•	•	1	1	2						
13						1	•	•	•	•	•	1	•	3					
14							1	•	•	•	•	•	1	1	1				
15							1	•	•	•	•	•	•	1	•	1			
16								1	•	•	•	•	•	•	1	1	3		
17								1	•	•	•	•	•	•	•	1	•	1	
18									1	•	•	•	•	•	•	•	1	1	2
19									1	•	•	•	•	•	•	•	•	1	•
20										1	•	•	•	•	•	•	•	•	1
21										1	•	•	•	•	•	•	•	•	•
22											1	•	•	•	•	•	•	•	•
23											1	•	•	•	•	•	•	•	•
24												1	•	•	•	•	•	•	•
25												1	•	•	•	•	•	•	•
26													1	•	•	•	•	•	•
27													1	•	•	•	•	•	•
28														1	•	•	•	•	•
29														1	•	•	•	•	•
30															1	•	•	•	•
31																1	•	•	•
32																	1	•	•
33																		1	•
34																			1
35																			1
36																			1
37																			1
38																			1
39																			1

Only the cases with an even number of vertices are shown. The odd cases consist of the complete graphs, with a single sporadic example which is four-regular with nine vertices. The members of the most interesting of the infinite sets are shown under the gray line. The sporadic examples are underlined in bold. The rest are members of the first four of the five infinite sets. The notation 2* means that only one of the two examples is sporadic.




5.4.4 The first three of the five infinite sets

These are very familiar, and already detailed in sections 5.3.6 to 5.3.8. They are the complete graphs, the complete bipartite graphs and the cocktail party graphs. Here are some additional comments on two of the three:

The complete graphs

The complete graphs K_n are $(n-1)$ -regular and integral Salem with spectrum $\{n-1, -1^{(n-1)}\}$. Despite having no vertices which are not adjacent, are usually regarded as trivially strongly regular.

K_3 is cyclotomic; the rest are Salem. K_4 can be drawn not only in circulant form and in planar/tetrahedral form, but also as a small Möbius ladder.




The cocktail party graphs

The cocktail party graphs CP_n are strongly regular with parameters $(n, n-2, n-4, n-2)$, and integral Salem with spectrum $\{n-2, 0^{(\frac{1}{2}n)}, -2^{(\frac{1}{2}n-1)}\}$.

Their adjacency matrices have first row with $a_0 = a_{\frac{1}{2}n-1} = 0$, and all other entries equal to one.

The smallest Salem cocktail party graph is CP_6 . Its adjacency matrix has first row $(0\ 1\ 1\ 0\ 1\ 1)$, and its spectrum is $\{4, 0^{(3)}, -2^{(2)}\}$. It is strongly regular with parameters $(6, 4, 2, 4)$, and is drawn both in circulant form and as a more familiar triangular anti-prism.



5.4.5 The fourth infinite set: the Bunga Bunga graphs¹

The cocktail party graph CP_n is the complete graph K_n with $\frac{1}{2}n$ independent edges removed. Similarly, the fourth infinite set consists of the complete bipartite graphs $K_{\frac{1}{2}n, \frac{1}{2}n}$ with $\frac{1}{2}n$ independent edges removed. The similarity is obvious and a related name is desirable. The independent edges removed from K_n are usually (not) drawn as those between opposite edges when the vertices form a regular polygon. Those removed from $K_{\frac{1}{2}n, \frac{1}{2}n}$ are usually thought of as the horizontal edges when the partitioned sets of vertices are drawn, as they often are, in parallel vertical columns.

The cocktail party graphs are so named because at a cocktail party everyone is said to talk to everyone else except their spouse. At a Bunga Bunga party (one imagines) everyone interacts with precisely those of the opposite gender who are not their partner, so the term “Bunga Bunga graph” seems appropriate.


It was noted above that $K_{\frac{1}{2}n, \frac{1}{2}n}$ has opposite diagonals only when $n = 2 \pmod{4}$, so they are there to be removed in only those cases. (0101010101) for example becomes (0101000101) but on the other hand (0101010101) cannot change in a similar way.

The Bunga Bunga graphs are $(\frac{1}{2}n-1)$ -regular with spectrum $\{\pm(n-1), \pm 1^{(n-1)}\}$. When $n = 0 \pmod{4}$ the graphs are still regular integral Salem, but not circulant.

BB_6 has the first row of its the adjacency matrix of $K_{3,3}$ is (010101) and to get $BB_{3,3}$ the bold underlined one has to be removed to leave (010001) which is clearly isomorphic to C_6 and cyclotomic.

BB_8 is isomorphic to the graph of the cube and is not circulant. When drawn bounded by a regular octagon, the “insides” lose the symmetry associated with circulant graphs.

BB_{10} is the first circulant Bunga Bunga Salem graph. (0101000101) is the first row of its adjacency matrix so the jumps are ± 1 and ± 3 and the graph decomposes into two interlocking decagons.



¹ The term “crown” graphs has been used, but is not widespread in the literature, has been applied to more than one type of graph, and does not reflect the close relationship to the cocktail party graphs.

5.4.6 The fifth infinite set

The complete bipartite graph $K_{n,n}$ has an even number of vertices and the first row of its adjacency matrix contains no even jumps and all possible odd jumps. This section is illustrated with the circulant Salem graphs which have 30 vertices, and the first row of the adjacency matrix of $K_{15,15}$ is

$$(01010101010101010101010101010101).$$

If precisely one pair of these odd jumps is removed, that which remains is a graph in the fifth infinite set. In each case these are subgraphs of the complete bipartite graph and so are still bipartite. In the case of 30 vertices this gives seven adjacency matrices with these first rows:

$$\begin{aligned} &(000101010101010101010101010100) \\ &(01000101010101010101010101010001) \\ &(01010001010101010101010101000101) \\ &(01010100010101010101010100010101) \\ &(01010101000101010101010001010101) \\ &(01010101010001010100010101010101) \\ &\text{and } (01010101010001000101010101010101). \end{aligned}$$

In cases where the structures generated by the missing jumps are isomorphic, the graphs which remain must also be isomorphic. This will happen when the highest common factor of the jump number and the number of vertices are the same in each case.

In the case of 30 vertices the jumps ± 1 , ± 7 , ± 11 and ± 13 each generate polygons with the full 30 vertices because $(30,1) = (30,7) = (30,11) = (30,13)$ so the matrices with these missing jumps generate isomorphic graphs, which in this case are trivial Salem with top eigenvalue 13. Similarly the jumps ± 3 and ± 9 each generate three decagons because $(30,3) = (30,9)$ so the matrices with these missing jumps also generate isomorphic graphs, which in this case are also trivial Salem with top eigenvalue 13 but with a different spectrum. Only the jump ± 5 generates five hexagons. No other jump number has the same HCF with 30 so there is only one circulant matrix representation of this graph. It is integral Salem with spectrum $\{\pm 13, \pm 2^{(4)}, \pm 1^{(10)}\}$.

The jumps of $(\pm)15$ give 15 copies of K_2 which are the opposite diagonals, and can be combined with each of the three or four 30-vertex polygons to form Möbius ladders.

So in total there are three non-isomorphic circulant Salem graphs with 30 vertices in the fifth infinite set. Besides these, and $K_{15,15}$ mentioned above, there are three other circulant Salem graphs with 30 vertices which are included here for completeness. They are the Bunga Bunga graph $BB_{15,15}$ which is $K_{15,15}$ with the diagonals removed, the cocktail party graph CP_{30} , and the complete graph K_{30} . The first rows of their adjacency matrices are:

$$\begin{aligned} &(01010101010101010001010101010101) \\ &(01111111111111110111111111111111) \\ &\text{and } (01111111111111111111111111111111). \end{aligned}$$


Members of the fifth infinite set with other numbers of vertices can be found in exactly the same way.

5.4.7 The seven sporadic circulant Salem graphs

The first graph is three-regular with six vertices. It is integral Salem with spectrum $\{3, 1, 0^{(2)}, -2^{(2)}\}$.

The first row of its adjacency matrix is (001110) so its circulant form decomposes into a pair of triangles with the diagonals of what would be the enclosing hexagon, as shown in the first drawing. This is our first example where there is no matrix of the preferred form $(01 \dots 1)$ which generates the same graph, so it cannot be drawn as a regular hexagon with circulant symmetrical “insides”. The second drawing, however, shows that it can be drawn as a regular hexagon with irregular insides.


The third drawing shows the more familiar form of a planar triangular prism.



The second graph is trivial Salem with spectrum $\{4, \sqrt{2}^{(2)}, 0, -\sqrt{2}^{(2)}, -2^{(2)}\}$. It is four-regular with eight vertices.


The first row of its adjacency matrix is (01100011) so its circulant form decomposes into an octagon and two squares as shown in the first drawing. It is perhaps best known as the quadrilateral antiprism.

A planar drawing is also shown.





This is the only circulant Salem graph with an odd number of vertices (nine) which is not complete.

It is four-regular and trivially Salem with top eigenvalue four. The first row of its adjacency matrix is (010011001) so the circulant form decomposes into two nonagons.



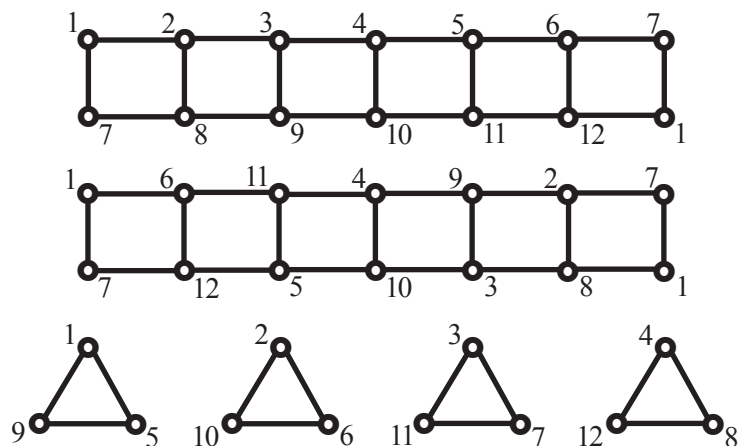
The fourth sporadic graph is integral Salem with spectrum $\{5, 2^{(2)}, 1^{(3)}, -1^{(6)}\}$. It is five-regular with twelve vertices.

This is another example of a circulant graph which cannot be represented by an adjacency matrix with its first row in the preferred form. The top row is (000110101100) so the circulant form decomposes into three squares, four triangles, and six copies of K_2 as shown below. These can now be assembled into a memorable representation of the graph. The six copies of K_2 are the diagonals of the squares which all go together to make three copies of K_4 . If the second square is now rotated through 90° and the third through 180° , the four triangles can be regarded as the edges between corresponding vertices of the copies of K_4 (i.e. top left to top left and so on). These are represented by thicker gray lines, each of which replaces four edges.



Similarly to the first sporadic example, it is possible to draw the vertices in a dodecagon with edges in the obvious places (for example $1-10-7-4-8-5-2-11-3-12-6-9-1$) but the “internal” symmetry would then be lost.

The fifth sporadic graph also has twelve vertices and is also integral Salem, with spectrum $\{7, 2^{(2)}, 1^{(2)}, -1, -2^{(6)}\}$, but it is seven-regular and has an adjacency matrix in the preferred form with first row (01001111001) . The jumps of ± 1 and ± 5 give two dodecagons, the jump of 6 adds the diagonals which turn them into Möbius ladders, and the jump of ± 4 gives four triangles.




The last two sporadic graphs both have 18 vertices.






The first is trivial Salem and 4-regular with largest eigenvalue four. The first row of its adjacency matrix is

(010000010001000001)

with jumps ± 1 and ± 7 , so it decomposes simply into two 18-vertex interlocking polygons.



The second is also trivial Salem and 5-regular with largest eigenvalue five. The first row of its adjacency matrix is (010100000100000101) with jumps ± 1 , ± 3 and 9, so it decomposes into a Möbius ladder and three hexagons.

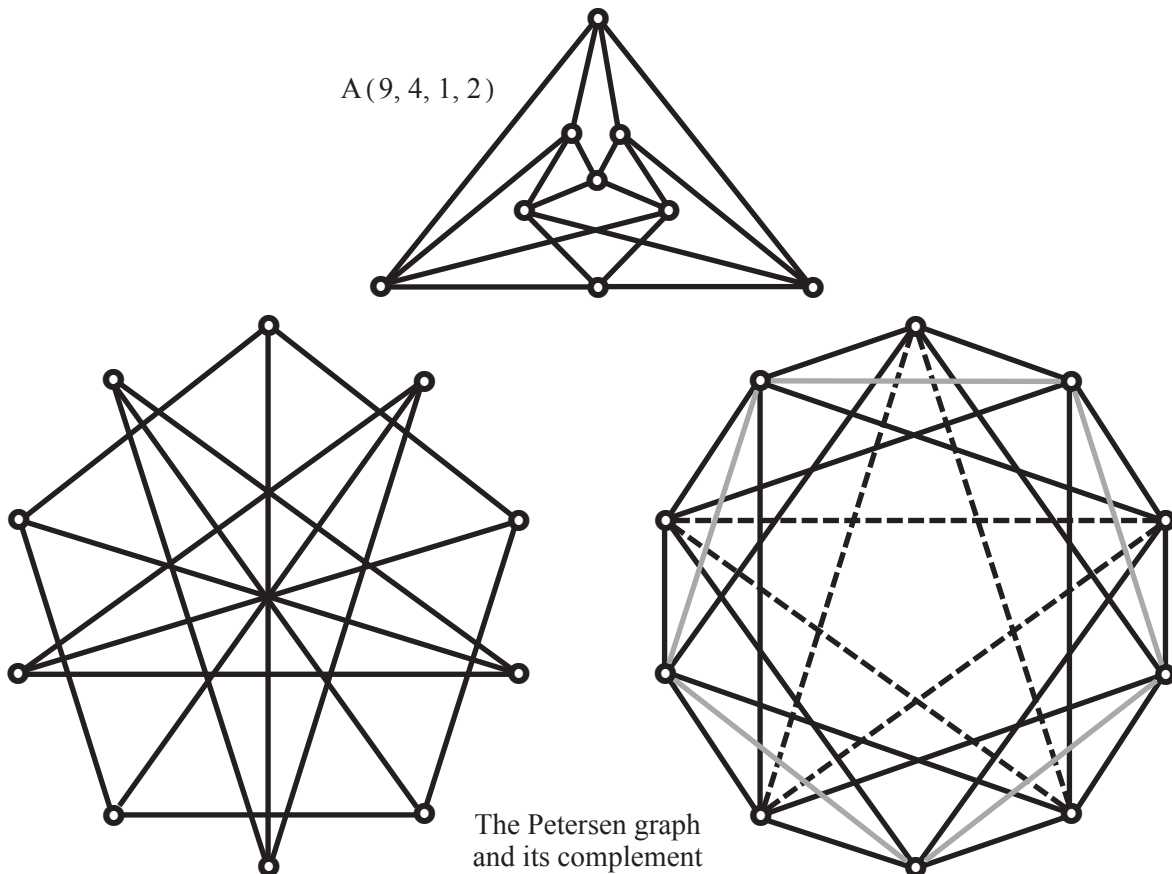


5.4.8 Small strongly regular Salem graphs

The table includes the complete graphs which are, at best, trivially strongly regular. Only three of the entries were not encountered in the previous section on circulant graphs. The parameters (v, r, a, n) mean that the graph has v vertices and is r -regular, with adjacent vertices sharing a neighbours and non-adjacent vertices sharing n neighbours.

Vertices	Total	Details
4	1	K_4
5	1	K_5
6	3	$K_{3,3}(6, 3, 0, 3)$ $CP_6(6, 4, 2, 4)$ K_6
7	1	K_7
8	3	$K_{4,4}(8, 4, 0, 4)$ $CP_8(8, 6, 4, 6)$ K_8
9	2	K_9 and $A(9, 4, 1, 2)$ with spectrum $\{4, 1^{(4)}, -2^{(4)}\}$ drawn below
10	5	$K_{5,5}(10, 5, 0, 5)$ $CP_{10}(10, 8, 6, 8)$ K_{10} The Petersen graph $(10, 3, 0, 1)$ with spectrum $\{3, 1^{(5)}, -2^{(4)}\}$ drawn below and Its complement $(10, 6, 3, 4)$ with spectrum $\{6, 1^{(4)}, -2^{(5)}\}$ also drawn below
11	1	K_{11}

Taking the usual drawing of the Petersen graph with an inner pentagram connected by spokes to an outer pentagon, the pentagram is here rotated through 180° and enlarged so that the vertices are now those of a regular decagon and the spokes become opposite diagonals. This unusual representation of the Petersen graph, not very interesting in itself, allows a reasonable attempt to be made at drawing a memorable version of its complement.




The Petersen graph
and its complement
92

5.4.9 Small regular Salem graphs


Table 5.5 The regular Salem graphs with fewer than twelve vertices arranged by their number of vertices and their n -regularity

n	Vertices							
	4	5	6	7	8	9	10	11
3	1		2		1		5	
4		1	1		3	3	1	
5			1		0	0	1	
6				1	1	0	1	
7					1	0	0	
8						1	1	
9							1	
10								1
Totals:	1	1	4	1	6	4	10	1


Six of these graphs are neither circulant nor strongly regular, so were not mentioned in the previous two sections. One is 3-regular with eight vertices, previously mentioned as the smallest non-circulant Bunga Bunga graph, isomorphic to the cube. It is also bipartite and appears in the next section. Graph A is one of those with eight vertices which are 4-regular. B is 4-regular with nine vertices. C, D and E are 3-regular with ten vertices. E is also bipartite and appears in the next section.




A with spectrum $\{4, 2, 0^{(3)}, -2^{(3)}\}$




B with spectrum $\{4, 2, 1^{(2)}, -1^{(2)}, -2^{(3)}\}$



C with spectrum $\{3, 2, 1^{(3)}, -1^{(2)}, -2^{(3)}\}$



D is trivial Salem






E with spectrum $\{\pm 3, \pm 2, \pm 1^{(2)}, 0^{(2)}\}$

5.4.10 Semiregular Salem graphs

This section is concerned with small connected Salem graphs which are bipartite, and in which all the vertices in the same partition are incident with the same number of edges. Bipartite graphs which are (fully) regular are included. There are 32 such graphs with fewer than twelve vertices. Of those, just seven are not complete bipartite graphs, and four of those seven are (fully) regular and have appeared in previous sections.

Vertices Total Details

5	1	$K_{3,2}$ (the star graph $K_{3,1}$ is cyclotomic)
6	3	$K_{3,3}$ $K_{4,2}$ $K_{5,1}$
7	3	$K_{4,3}$ $K_{5,2}$ $K_{6,1}$
8	5	$K_{4,4}$ $K_{5,3}$ $K_{6,2}$ $K_{7,1}$ and BB_8 which is isomorphic to the cube
9	5	$K_{5,4}$ $K_{6,3}$ $K_{7,2}$ $K_{8,1}$ and graph $A(9,4,2)$ shown below, a subgraph of $K_{6,3}$
10	10	$K_{5,5}$ $K_{6,4}$ $K_{7,3}$ $K_{8,2}$ $K_{9,1}$ The three-regular bipartite graph shown in the previous section The first graph in the fifth infinite set of circulant Selem graphs BB_{10} , the first circulant Bunga Bunga graph Graphs $B(10,3,2)$ and $C(10,3,2)$ shown below, subgraphs of $K_{6,4}$
11	5	$K_{6,5}$ $K_{7,4}$ $K_{8,3}$ $K_{9,2}$ $K_{10,1}$



Chapter 6

Non-bipartite graphs of small Mahler measure

6.1 Introduction

6.1.1 Publication

This is an exposition of the paper (with the same title as this chapter) published in the Journal of Combinatorics and Number Theory (ISSN 1942–5600), Volume 5 (2013), Number 2, pp. 53–64. The contributors were (in alphabetical order) the present author and James McKee of Royal Holloway, University of London, and Chris Smyth of the University of Edinburgh.

6.1.2 Notes on the title

Bipartite graphs of small Mahler measure were classified in a previous paper [MS] by the second and third authors, so this paper completes the classification of *all* graphs with small Mahler measure.

The word “small” here means less than the Golden Ratio $\phi := \frac{1+\sqrt{5}}{2}$. This is shown to be the least limit point of the set of Mahler measures of non-bipartite graphs, so it is the natural choice.

It is worth repeating that for Salem graphs the associated Salem number and the Mahler measure are exactly the same.

6.1.3 Contributions of the present author

The main result in the paper is that every non-bipartite graph with small Mahler measure belongs either to one of three infinite families or to a set of eight sporadic examples. One of the families (the odd cyclic graphs) has been known since 1970 [Sm] to consist of the only connected non-bipartite graphs which are cyclotomic (i.e. with Mahler measure equal to one). The present author completed the classification, providing all remaining discoveries, calculations, conjectures and motivation for the paper.

6.1.4 Differences between this exposition and the published paper

- The introduction has been completely rewritten since the context here is very different from that of an isolated paper. In particular, ideas which have been carefully introduced (or at least, it is hoped, adequately mentioned) in chapter 1, have not been reintroduced here. Similarly Cauchy’s interlace theorem, and its crucial importance in this thesis, has already been carefully presented in chapter 3.
- The diagram on page 97 has been added to show more clearly how graphs in the classification can be systematically developed in various ways from C_3 (the ‘fundamental’ non-bipartite graph), since this was a major contribution by the present author.

- Because space is not a major consideration here, a little more explanation and comment are included.
- Non-connected graphs are not discussed in the rest of this thesis because nothing essentially new was found that was worth saying about them. The same could be said here, but, because the paper complements [MS] where they were explicitly considered, they are fully treated, especially in Lemma 6.2.1.

6.1.5 The bipartite case

Bipartite graphs with small Mahler measure were classified in [MS]. This section outlines the method.

Recall from the first chapter that if a graph G has characteristic polynomial $\chi_G(x)$ then the Mahler measure of the graph $M(G)$ can be written in terms of the eigenvalues of G :

$$M(G) = \prod_{\chi_G(x)=0, |\lambda|>2} \frac{|\lambda| + \sqrt{\lambda^2 - 4}}{2}.$$

Let $\theta = \sqrt{2 + \sqrt{5}} = 2.058\dots$. If $\lambda = \theta$ then $\frac{|\lambda| + \sqrt{\lambda^2 - 4}}{2} = \sqrt{\phi}$.

But if a graph is bipartite then its roots are symmetric about the origin [CR], so λ and $-\lambda$ both contribute to its Mahler measure, and if its Mahler measure is less than ϕ then its spectral radius must be less than θ .

All connected graphs with largest eigenvalue in the interval $(2, \theta]$ have been described by Cvetković and Rowlinson [CvRo Theorem 2.4] drawing on [BN] and [CDG]. So the work of [MS] was to find which of these graphs have Mahler measure below ϕ .

In the non-bipartite case, however, it is possible for the spectral radius to be larger, with the Mahler measure still below ϕ , so a different approach is required.

6.2 Theorem

Every connected non-bipartite graph with Mahler measure less than ϕ is one of the following:

- an odd cycle
- a ‘kite’ graph (Figure 6.1)
- a ‘balloon’ graph (Figure 6.2)
- one of eight sporadic examples (Figure 6.3)

The diagram on the next page shows the interrelationships between them.




Figure 6.1. The kite graphs: Kt_n has n vertices




Figure 6.2. The balloon graphs: Bl_{2n} has $2n$ vertices





Figure 6.3. The sporadic graphs: Sp_a to Sp_h



6.2.1 Nonconnected graphs

By using [MS, Theorem 10.2] and the following remark we can extend theorem 6.2 to include non-connected graphs.

Corollary Every non-bipartite graph of small Mahler measure is one of the following:

- a (not necessarily connected) bipartite graph of small Mahler measure with one or more additional components consisting of odd cycles;
- a graph with one connected component as given in Theorem 6.2 with any other components cyclotomic
- a graph with one connected component Bl_8 , one connected component the tree



and any other components cyclotomic.

6.2.2 A lower bound

As an immediate consequence of Theorem 6.2 and the computations involved in its proof we find the following lower bound on Mahler measures greater than 1 for connected non-bipartite graphs.

Corollary Let G be a connected non-bipartite graph. Then either $M(G) = 1$ or

$$M(G) \geq M(Bl_8) = 1.35098\dots, \text{ the larger real root of } z^{10} - z^9 - z^6 + z^5 - z^4 - z + 1.$$

The polynomial is obtained by substituting $n=4$ in expression (\star) in §6.4, removing the first factor which is just the cyclotomic polynomial often known as Φ_7 , and carrying out the division.

6.2.3 Plan of the proof

After Smith's result [Sm] we need only consider non-cyclotomic graphs.

- In 6.3 we prove that all kites have small Mahler measure;
- In 6.4 we prove that all balloons have small Mahler measure;
- In 6.5 we describe computations which deal with all small examples;
- In 6.6 we list some graphs with Mahler measure which is not small; by interlacing they cannot be induced subgraphs of graphs with small Mahler measure;
- In Lemma 6.7.1 we prove that any connected graph which has small Mahler measure and contains a triangle must be a kite; and finally
- In Lemma 6.7.3 we show that all remaining connected non-bipartite graphs with small Mahler measure are balloons.

We finally prove Corollary 6.2.1.

6.3 All kites have small Mahler measure

Kt_n is a line graph [GR, §1.7], so has all eigenvalues in the interval $[-2, \infty)$ [B, Proposition 3.7]. Deleting one of the vertices in the triangle leaves a cyclotomic graph as is seen from Smith's classification [S]. But Kt_n itself is not one of Smith's graphs, so does not have all its eigenvalues in $[-2, 2]$, so by interlacing it is a Salem graph with a unique eigenvalue larger than two which is the only eigenvalue that contributes to the Mahler measure.

Let λ_n be the largest eigenvalue of Kt_n . As n increases so does λ_n , and indeed it strictly increases [GR, Theorem 8.8.1(b)]. Write $\lambda_n = z_n + \frac{1}{z_n}$, with $z_n > 1$; then z_n also strictly increases with n , and equals the Mahler measure of Kt_n . By [MS, Lemma 4.3], using the explicit formula in the proof of [MS, Lemma 4.1], z_n converges to a root of $z^2 - z - 1 = 0$, and it must be the positive root ϕ .

Hence $z_n = M(Kt_n) < \phi$ for all $n \geq 4$, and we see that ϕ is a limit point of the set of Mahler measures of non-bipartite graphs.

Using Lemma 4.1 of [MS], we compute that the reciprocal polynomial of Kt_n is

$$z^{2n} - 2z^3 + 1 - z^4(2z+1)\frac{z^{2n-6}-1}{z^2-1}.$$

Table 6.1 shows the first few values of $M(Kt_n)$.

Table 6.1. Mahler measures of small kites and balloons

n	$M(Kt_n)$	$M(Bl_{2n})$
2		1.5061...
3		1.4012...
4	1.5061...	1.3509...
5	1.5823...	1.5064...
6	1.6054...	1.5783...
7	1.6134...	1.6020...
8	1.6162...	1.6113...
9	1.6173...	1.6151...
10	1.6177...	1.6168...

6.4 All balloons have small Mahler measure

Balloons cause more trouble than kites because, apart from small cases, they have two eigenvalues outside the interval $[-2, 2]$. As indicated in Table 6.1, the Mahler measures of balloons initially decrease as the number of vertices grows, reaching a minimum for Bl_8 , then appear to increase towards ϕ . This we now prove.

Computing the characteristic polynomial by expanding along the row corresponding to the leaf, the reciprocal polynomial of Bl_{2n} is found to be

$$\frac{z^{2n-1}-1}{z-1} \cdot \frac{(z^4-z^2-1)z^{2n-1}-(z^4+z^2-1)}{z+1}. \quad (\star)$$

Since we are interested in the Mahler measure we can delete the left factor which is cyclotomic. It also proves easier to multiply by $z+1$ and work with the numerator of the right factor, rather than working with the polynomial which results from the division. This also does not change the Mahler measure.

$$(z^4-z^2-1)z^{2n-1}-(z^4+z^2-1) = P_n(z) \text{ say.}$$

To show that Bl_{2n} has small Mahler measure we must show that $M(P_n) < \phi$. For $n < 5$ we check this by direct computation. It remains to deal with $n \geq 5$.

Deleting the vertex of valency three leaves a (disconnected) cyclotomic graph so, by interlacing, P_n has at most two roots outside the unit disc. Note that

$$P_n(-\sqrt{\phi}) < 0 \quad P_n(-1) = 0 \quad P_n'(-1) = 9 - 2n < 0 \text{ for } n \geq 5 \quad P_n(\sqrt{\phi}) < 0 \quad \text{and} \quad P_n(\infty) = +\infty$$

so that for $n \geq 5$, P_n has a root z_n^- in $(-\sqrt{\phi}, -1)$ and a root z_n^+ in $(\sqrt{\phi}, \infty)$, and these account for all possible roots outside the unit disc.

From $P_n(z) = 0$ we get

$$\frac{\log \left| \frac{z^4+z^2-1}{z^4-z^2-1} \right|}{\log |z^2|} = \frac{2n-1}{2}.$$

Putting $z^2 = (1+x)\phi$ and remembering that $\phi - 1 - \frac{1}{\phi} = 0$ and $2\phi - 1 = \phi + \frac{1}{\phi}$, the left-hand side becomes

$$g(x) := \frac{\log \left| 1 + \frac{2(1+x)}{x(\phi x + \phi + \frac{1}{\phi})} \right|}{\log \phi + \log |1+x|}.$$

The two roots z_n^- and z_n^+ correspond to real roots of the equation $g(x) = \frac{2n-1}{2}$. Call these $-u_n'$ and u_n where $z_n^- = \sqrt{(1-u_n')\phi}$ and $z_n^+ = \sqrt{(1+u_n)\phi}$.

$g(x)$ is decreasing for $x > 0$ because its numerator is decreasing and its denominator is increasing. Since $g(0.1) < \frac{9}{2}$ we see from 0.0 that $u_n \in (0, 0.1)$ for all $n \geq 5$.

We have

$$M(P_n) = \phi \sqrt{(1+u_n)(1-u_n')} = \phi \sqrt{(1+u_n-u_n'-u_n u_n')},$$

which is less than ϕ if $u_n' > u_n$. We now show that this is indeed the case.

Knowing that $g(x)$ is decreasing in $(0, \infty)$ and $u_n \in (0, 0.1)$, it will be enough to show that

$$g(-x) > g(x) \text{ for } x \in (0, 0.1).$$

On rewriting the numerator of $g(x)$ as

$$\log\left|\frac{1}{x}\right| + C + \log|1+xR(x)|$$

where

$$C = \log\frac{2\phi}{\phi+2} \approx -0.11157 \quad \text{and} \quad R(x) = \frac{(4\phi+3)x+7\phi+5}{(4\phi+2)x+6\phi+2},$$

one readily checks that

$$g(-x) - g(x) \sim \frac{2x}{(\log\phi)^2} \log\left|\frac{1}{x}\right|$$

as $x \rightarrow 0^+$, and using the error term in the Taylor approximation to the logarithm function shows that this positive main term dominates in the interval $(0, 0.1)$ as desired.

6.5 Details of some computations

The connected graphs with small Mahler measure can be grown in exactly the same way as the small Salem graphs were grown in chapter 3. But in the present case the process can be made very much more efficient by using the fact that non-bipartite graphs are precisely those which contain at least one odd cycle.

We start with a triangle and grow all graphs with four and then five vertices, at each step keeping only those with small Mahler measure and only one representative of each isomorphism class. At this stage we add in ‘manually’ the 5-cycle, which cannot be grown from a triangle. We can proceed in this way for as long as computational limitations allow. This gives the following lemma:

6.5.1 Lemma. Let G be a connected non-bipartite graph with $1 < M(G) < \phi$ and with at most eight vertices. Then G is either a kite (Figure 6.1), a balloon (Figure 6.2) or one of $Sp_a, Sp_b, Sp_d, Sp_e, Sp_h$ (Figure 6.3).

These results are shown in Table 6.2.

Table 6.2 Small connected graphs with Mahler measure strictly between 1 and ϕ

Number of vertices	Graph(s)	Mahler measure(s)
4	$Kt_4 = Bl_4$	1.5061...
5	Kt_5	1.5823...
6	Kt_6, Bl_6	1.6054... 1.4012
7	Kt_7, Sp_a, Sp_d, Sp_h	1.6134... 1.4723... 1.5560... 1.5823...
8	Kt_8, Bl_8, Sp_b, Sp_e	1.6162... 1.3509... 1.4967... 1.5991...

Once the proof of theorem 6.2 is complete, the table can be extended at will. There are three more sporadic examples (Sp_c, Sp_f, Sp_g), all with nine vertices, and beyond that there are only kites and balloons.

For the proof we need only grow up to eight vertices as later lemmas deal with all graphs on nine or more. It is striking to notice, however, that once we have grown to ten vertices there are no more graphs containing pentagons or heptagons; the growing process has terminated in all possible directions so there are only a finite number. In particular:

6.5.2. Lemma. The only connected, non-bipartite graphs with Mahler measure in the interval $(1, \phi)$ which contain either a 5-cycle or a 7-cycle are the balloons Bl_6 and Bl_8 , and the eight sporadic graphs of Figure 3.

6.6. Some graphs that do not have small Mahler measure

These are required later in the proof. The first four are bipartite, so we can appeal to [MS Theorem 10.2].

6.6.1. Lemma. The four graphs L_1, L_2, L_3, L_4 in Figure 6.4 all have Mahler measure greater than ϕ .

The proof is by direct computation.




Figure 6.4. The graphs L_1 to L_4




Figure 6.5. The tailed kites $\tilde{K}t_n$ with $n \geq 5$




Figure 6.6. The graphs $Q(a, b, c)$ and $\tilde{Q}(d, e)$

6.6.2. Lemma. The ‘tailed kites’ $\tilde{\text{Kt}}_n$ of Figure 5 (n vertices, $n \geq 5$) all have Mahler measure greater than ϕ .

Proof. Note that $\tilde{\text{Kt}}_n$ is not one of Smith’s graphs [S], but that deleting one of the degree-2 vertices in the triangle of $\tilde{\text{Kt}}_n$ leaves a subgraph of one of those graphs. By interlacing, $\tilde{\text{Kt}}_n$ has at least one eigenvalue greater than 2, and indeed exactly one since the spectral radius of a graph is always equal to one of its eigenvalues [GR, Lemma 8.7.3]. On the other hand, $\tilde{\text{Kt}}_n$ is a generalised line graph [B, 3h], so has all eigenvalues at least -2 . Thus $\tilde{\text{Kt}}_n$ has a unique eigenvalue outside the interval $[-2, 2]$, and this is >2 . From [HS, Proposition 2.4], the Mahler measure of $\tilde{\text{Kt}}_n$ slightly decreases as n increases. In the limit using [MS, §4], this sequence of Mahler measures converges to ϕ . Hence $M(\tilde{\text{Kt}}_n) > \phi$ for all $n \geq 5$. \square

6.6.3. Lemma. Let $\tilde{Q}(d, e)$ be the graph shown in Figure 6, where $d, e \geq 1$ and $d + e > 2$. Then, with the exceptions of $\text{Sp}_d, \text{Sp}_g, \text{Sp}_h$, (corresponding to $(d, e) = (2, 3), (3, 4), (1, 4)$) one has $M(\tilde{Q}(d, e)) > \phi$.

Proof. We may assume that $d \leq e$. For $e < 9$ we check the result by direct computation. for $e \geq 9$, delete suitable vertices from the middle of the longer path between the two degree-3 vertices to leave a subgraph $Q(a, b, c)$ (see Figure 6; here $(a-1) + (c-1) \leq 7 < e-1$) in the following list:

$$Q(3, 1, 3), Q(3, 2, 3), Q(3, 3, 3), Q(3, 4, 4), Q(3, 5, 5), Q(3, 6, 5), Q(3, 7, 6), Q(4, 8, 5)$$

or $Q(4, d, 4)$ if $d \geq 9$.

From the computations in the proof of [MS, Theorem 10.2], this (bipartite) subgraph has Mahler measure greater than ϕ , and hence by interlacing so does $\tilde{Q}(d, e)$. \square

6.7. All large enough, connected, non-cyclotomic, non-bipartite graphs of small Mahler measure are either kites or balloons

6.7.1 Lemma. Let G be a connected graph with Mahler measure in the interval $(1, \phi)$. If G contains a triangle, then G is a kite.

Proof. We use induction on $n \geq 1$. For $n \leq 8$ the direct computations in §6.5 establish the result.

Suppose that $n > 8$ and that the result is known for relevant graphs with fewer vertices. Let T be a triangle in G and for any vertex v define the distance from v to T to be the minimal number of edges in a path from v to one of the vertices in T . Take v a vertex of maximal distance from T . Let G' be the subgraph obtained by deleting v together with all its incident edges. Maximality of the distance from v to T ensures that G' is connected. By interlacing, the Mahler measure of G' is at most that of G , so either equals 1 or is in the interval $(1, \phi)$. The former is excluded by inspection of Smith’s graphs [S], so by our inductive hypothesis $G' = \text{Kt}_{n-1}$. Let x be the leaf in G' , with y its neighbour. By maximality of the distance of v from T , the only possible neighbours of v in G are x and y .

First consider the possibility that v is adjacent to both x and y . Using $n-1 \geq 8$, we could then delete vertices from the middle of the path from y to T to leave two disjoint copies of Kt_4 . By interlacing we would have $M(G) \geq M(Kt_4)^2 > 1.50613^2 > \phi$, contradicting $M(G) < \phi$. We deduce that v is adjacent to exactly one of x and y .

Next consider the possibility that v is adjacent to y only. Then G is a tailed kite (Figure 5) and lemma 6.6.2 gives a contradiction.

We are forced to the conclusion that v is adjacent to x only, and therefore that $G = Kt_n$. □

6.7.2 Lemma. Let G be a connected, non-bipartite graph, with Mahler measure in the interval $(1, \phi)$. Let C be an odd cycle in G of shortest length. If v is a vertex not in C , then v is adjacent to at most one vertex of C .

Proof. If G contains a triangle, then the result follows from Lemma 6.7.1. We may therefore suppose that G contains no triangles.

Suppose that v is a vertex not in C that is adjacent to two vertices x and y on C (and perhaps adjacent to others). The cycle provides us with two paths from x to y , and since C has odd length one of these paths P contains an even number of edges. If P had more than two edges, then following the odd-length path from x to y , then going from y to v and from v to x would give an odd cycle shorter than C . Hence P has exactly two edges; let z be the vertex on P between x and y , and let u be the other neighbour of y on C . Since G has no triangles, and u cannot be a neighbour of x (else we could shorten C by replacing the path $xzyu$ by the path xu) the subgraph included by x, y, z, u, v is L_3 in Figure 4. Lemma 6.6.1 records that $M(L_3) > \phi$, hence by interlacing we have $M(G) > \phi$ which is a contradiction. We conclude that no such vertex v exists, which is the claim of the current Lemma. □

6.7.3. Lemma. Let G be a connected non-bipartite graph with Mahler measure in the interval $(1, \phi)$. Suppose that G has n vertices and that the shortest odd cycle in G has length $2m-1$. If $m \geq 5$ then $G = Bl_{2m}$.

Proof. We use induction on n . For $n \leq 9$ the result is vacuous.

We suppose that $n > 9$ and that the result is known for all relevant smaller graphs. Let C be a shortest odd cycle in G . We may assume that C has at least 9 edges, or there is nothing to prove. Since $M(C) = 1$ there must be other vertices in G . Let v be a vertex in G that is as far distant from C as possible. Deleting v leaves a connected graph H , containing C as its shortest odd cycle. If $M(H) = 1$ then $H = C$ (Case 1). Otherwise, by our inductive hypothesis, $n-1$ is even and $H = Bl_{n-1}$ (Case 2); we now show that this case cannot arise.

Case 1: $H=C$. Then $n-1$ is odd so n is even, and by Lemma 6.7.2 $G=Bl_n$.

Case 2: $H=Bl_{n-1}$. Let x be the leaf of H , and let y be its neighbour on C . We consider three subcases:

- 2a: v is adjacent to x only;
 - 2b: v is adjacent to x and to a vertex z on C (exactly one such neighbour on C , after Lemma 6.7.2);
 - 2c: v is adjacent to a vertex z on C (again unique after Lemma 6.7.2), but not to x .
- 2a: Noting that L_1 of Figure 4 is an induced subgraph, we see that this case is ruled out by Lemma 6.6.1.
 - 2b: Consider the path P on C that connects y and z by an odd number of edges. By minimality of the length of C , the only possible lengths for P are 1 and 3 (otherwise we could find a shorter odd closed walk by replacing the path P within C by the path zvx). If P has length 1, then G contains L_3 of Figure 4 as an induced subgraph; if P has length 3 then it contains L_4 . In either case we see that Lemma 6.6.1 gives a contradiction.
 - 2c: This gives two further subcases. If $z=y$, then we have L_2 of Figure 4 as an induced subgraph of G . If $z \neq y$, then we appeal to Lemma 6.6.3, noting that $m \geq 5$ excludes the sporadic cases.

Each subcase of Case 2 produces a contradiction, so we have Case 1: $H=C$ and $G=Bl_n$. \square

This completes the proof of the main Theorem 6.2. \square

6.8 Proof of Corollary 6.2.1

The proof follows from Theorem 6.2 using the facts that a graph is non-bipartite if and only if at least one of its components is non-bipartite, and that the Mahler measure of a graph is the product of the Mahler measures of its connected components.

Let G be a non-cyclotomic graph of small Mahler measure.

- If all the non-cyclotomic components are bipartite then at least one cyclotomic component must be non-bipartite, and so an odd cycle. This gives the first case.
- Otherwise G has a non-bipartite non-cyclotomic component as described by the Theorem. As all of these have Mahler measure at least $M(Bl_8)=1.350980338 > \sqrt{\phi}$, there can be only one of these components. If all the other components are cyclotomic, we have the second case.

- Otherwise some component is non-cyclotomic and bipartite, in which case, by [MS, Theorem 10.2], it has Mahler measure at least $M(T)=1.76280818$, where T is the tree in the Corollary. But then the non-bipartite component of G can have Mahler measure at most $\frac{\phi}{1.76280818} = 1.375550773$. But Bl_8 is the only such non-bipartite non-cyclotomic connected graph, all others having Mahler measure at least $M(\text{Bl}_6)=1.401268368$, and the only connected bipartite non-cyclotomic graph that has Mahler measure below $\frac{\phi}{M(\text{Bl}_8)}$ is T . This gives the third case.

6.9 Final remarks

The paper ends by pointing out that very little is known about graphs of slightly less small Mahler measure (beyond the ϕ boundary), and by suggesting that the work could be extended to include, for example, signed graphs or, more generally, integer symmetric matrices.

Bibliography

- [A] F S Acton, *Numerical Methods That Work*, Washington, DC: Math. Assoc. Amer., 1990, 334.
- [BDKM] D Bindel, J Demmel, W Kahan, O Marques, On computing Givens rotations reliably and efficiently, *LAPACK Working Note 148*, University of Tennessee, 2001.
- [BEG] H Bass, D R Estes, R M Guralnick, Eigenvalues of Symmetric Matrices and Graphs, *Journal of Algebra* 168, 536–567 (1994).
- [Bi] N Biggs, *Algebraic Graph Theory*, second edition, Cambridge University Press, Cambridge, 1993.
- [Bo] P Borwein, *Computational Excursions in Analysis and Number Theory*, Springer-Verlag, 2002, 16.
- [BN] A E Brouwer, A Neumaier, The graphs of spectral radius between 2 and $\sqrt{2+\sqrt{5}}$, *Linear Algebra Appl.* 114/115 (1989), 273–276.
- [CR] C A Coulson, G S Rushbrooke, Note on the method of molecular orbitals, *Math. Proc. Camb. Philos. Soc.* 36 (1940), 193–200.
- [CDG] D Cvetković, M Doob, I Gutman, On graphs whose spectral radius does not exceed $\sqrt{2+\sqrt{5}}$, *Ars Combin.* 14 (1982), 225–239.
- [CvRo] D Cvetković, P Rowlinson, The largest eigenvalue of a graph: a survey, *Linear and Multilinear Algebra* 28 (1990), 3–33.
- [ERS] P Erdős, A Rényi, V Sós, On a problem of graph theory, *Studia Sci. Math. Hungar.* 1966, 1: 215–235
- [FSG] P Foggia, S Sansone, M Vento, A Performance Comparison of Five Algorithms for Graph Isomorphism, *Proc. of the 3rd IAPR-TC-15 International Workshop of Graph-based Representations in Pattern Recognition* (2001), 188–199
- [GM] L Gumbrell, J F McKee, A classification of all 1-Salem graphs, *LMS Journal of Computation and Mathematics* 17 (2014), 582–594.
- [GR] C Godsil, G Royle, *Algebraic Graph Theory*, Grad. Texts in Math. 207, Springer, New York, 2001.
- [H] Suk-Guen Hwang, Cauchy’s Interlace Theorem for Eigenvalues of Hermitian Matrices, *The Mathematical Association of America* (2004), 111, 157-158.
- [HS] A J Hoffman, J H Smith, On the spectral radii of topologically equivalent graphs, *Recent Advances in Graph Theory (Proc. Second Czechoslovak Sympos., Prague, 1974)*, edited by M Fielder, pp. 273–281, Prague: Academia, 1975.

- [M] J F McKee, Conjugate Algebraic Numbers on Conics: A Survey, *London Mathematical Society Lecture Note Series* 352 (2008) 212–214.
- [MS] J F McKee, C J Smyth, Salem Numbers, Pisot numbers, Mahler measure, and graphs, *Exp. Math.* 14 (2005), 211–229.
- [P] The PARI Group, version 2.3.4, Bordeaux 2006, available from <http://pari.math.u-bordeaux.fr/>.
- [Sh] S S Shrikhande, The uniqueness of the L2 association scheme, *Annals of Mathematical Statistics*, 1959 30: 781–798.
- [Sm] J H Smith, Some properties of the spectrum of a graph, in: *Combinatorial Structures and their Applications (Proc. Calgary Internat. Cong., Calgary, Alta., 1969)*, edited by R Guy, H Hanani, H Saver, J Schönheim, pp. 403–406, Gordon and Breach, New York, 1970.
- [W1] E W Weisstein, *CRC Concise Encyclopedia of Mathematics*, Chapman and Hall CRC, 2002, 646–649.
- [W2] E W Weisstein, *CRC Concise Encyclopedia of Mathematics*, Chapman and Hall CRC, 2002, 855–857.